

AD-R123 464

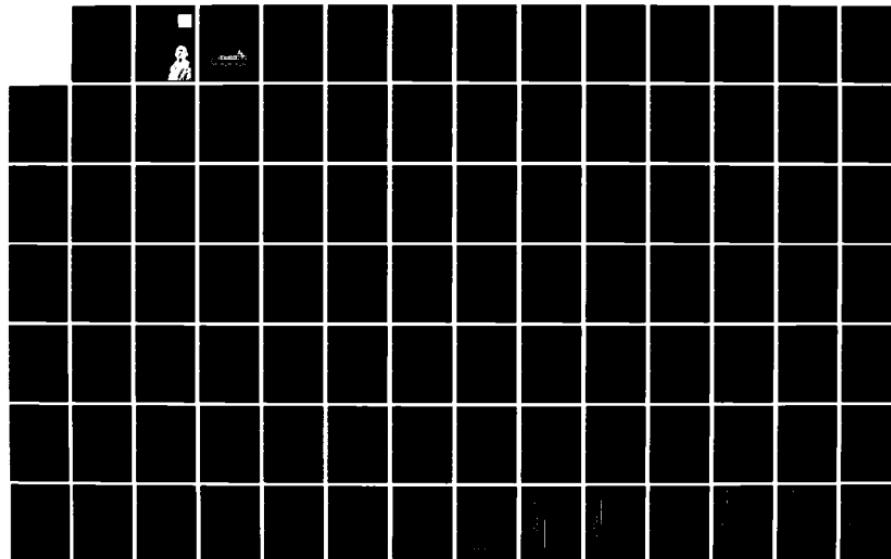
PROCEEDINGS OF A SYMPOSIUM ON CARGO SHIP ROUTING AND
SCHEDULING HELD WASH. (U) GEORGE WASHINGTON UNIV
WASHINGTON DC INST FOR MANAGEMENT SCIE. R M SOLAND

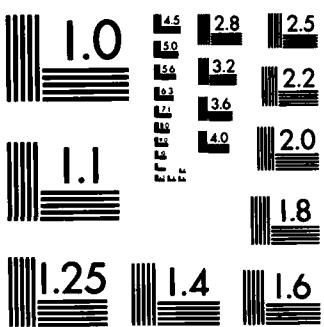
1/4

UNCLASSIFIED

15 DEC 82 SERIAL-TM-69250 N00014-82-G-0016 F/G 15/5

NL



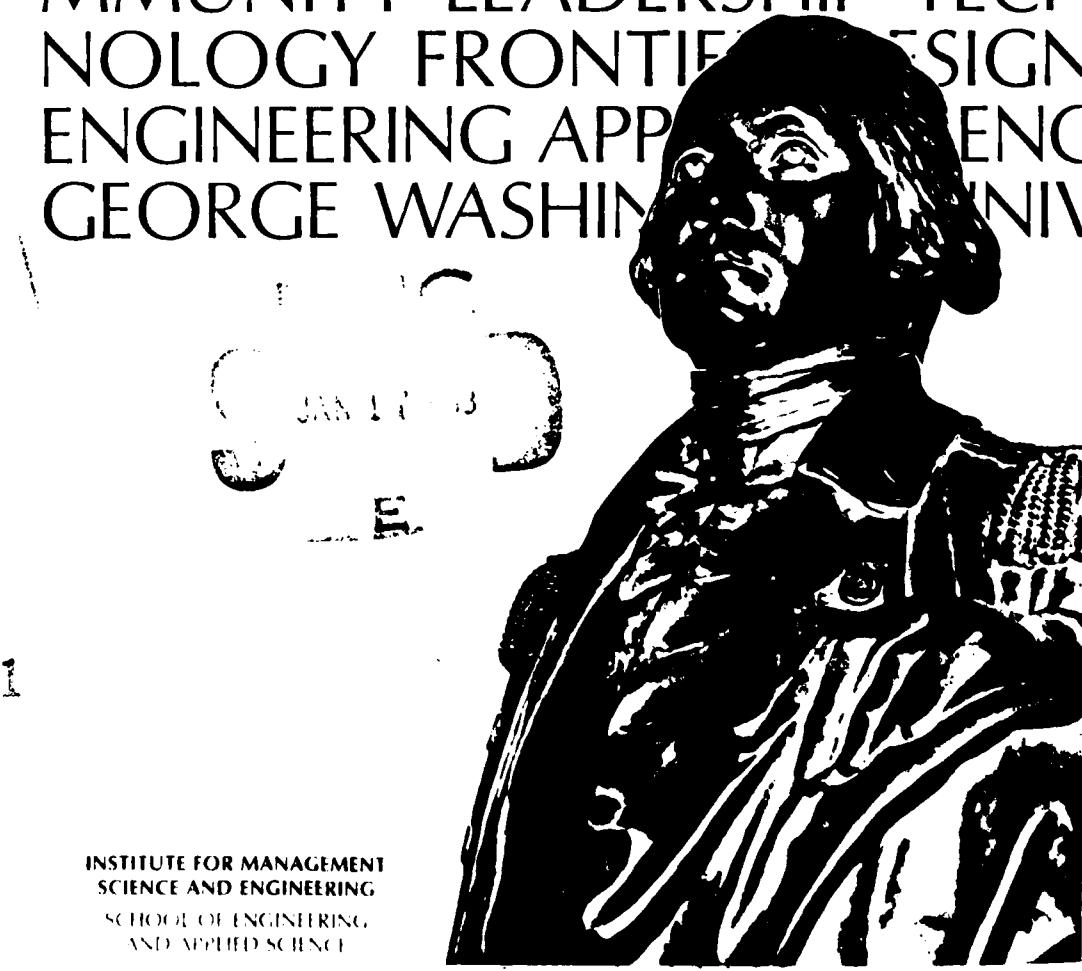


MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(12)

THE
GEORGE
WASHINGTON
UNIVERSITY

STUDENTS FACULTY STUDY R
ESEARCH DEVELOPMENT FUT
URE CAREER CREATIVITY CC
MMUNITY LEADERSHIP TECH
NOLOGY FRONTIERS DESIGN
ENGINEERING APPRENTICESHIP
GEORGE WASHINGTON UNIVERSITY



03 01

THIS FILE COPY

DMIS

INSTITUTE FOR MANAGEMENT
SCIENCE AND ENGINEERING
SCHOOL OF ENGINEERING
AND APPLIED SCIENCE

THE GEORGE WASHINGTON UNIVERSITY
School of Engineering and Applied Science
Institute for Management Science and Engineering

PROCEEDINGS OF A SYMPOSIUM ON
CARGO SHIP ROUTING AND SCHEDULING

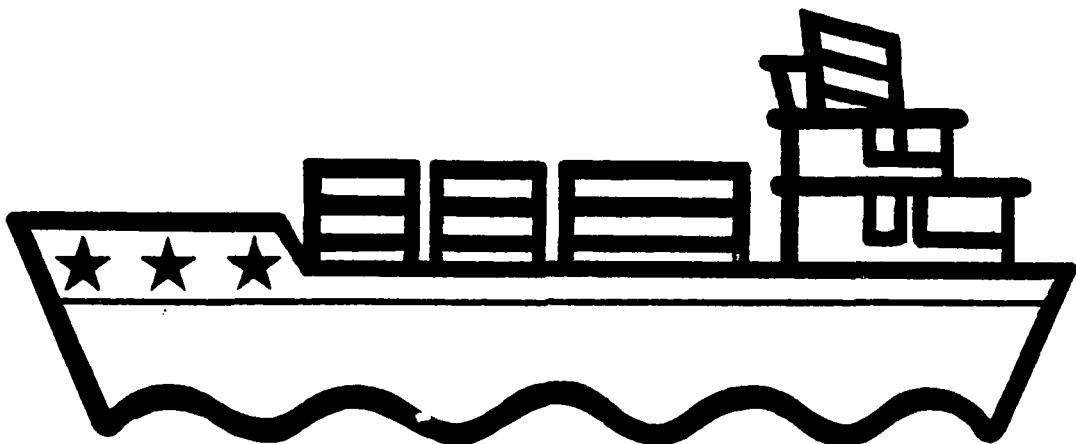
Washington, DC, February 3, 4, 1982

Edited by

Richard M. Soland
Department of Operations Research
School of Engineering and Applied Science
The George Washington University

TECHNICAL MEMORANDUM

Serial TM-69250
December 15, 1982



Grant N00014-82-G-0016
Office of Naval Research

This document contains neither recommendations nor conclusions of the Defense Science Board. It is the property of the Defense Science Board and is loaned to your agency; it is to be returned upon request.	Approved
---	----------

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER TM-69250	2. GOVT ACCESSION NO. A123 464	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) PROCEEDINGS OF A SYMPOSIUM ON CARGO SHIP ROUTING AND SCHEDULING	5. TYPE OF REPORT & PERIOD COVERED FINAL	
7. AUTHOR(s) EDITED BY RICHARD M. SOLAND	6. PERFORMING ORG. REPORT NUMBER TM- 69250	
9. PERFORMING ORGANIZATION NAME AND ADDRESS THE GEORGE WASHINGTON UNIVERSITY SCHOOL OF ENGINEERING AND APPLIED SCIENCE WASHINGTON, DC 20052	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS OFFICE OF NAVAL RESEARCH CODE 614A-PAE ARLINGTON, VIRGINIA 22217	12. REPORT DATE 15 December 1982	
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)	13. NUMBER OF PAGES 347	
15. SECURITY CLASS. (of this report) UNCLASSIFIED		
15a. DECLASSIFICATION/DOWNGRADING SCHEDULE		
16. DISTRIBUTION STATEMENT (of this Report) APPROVED FOR PUBLIC SALE AND DISTRIBUTION; DISTRIBUTION UNLIMITED.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) SCHEDULING ROUTING CARGO SHIPS SYMPOSIUM		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A Symposium on Cargo Ship Routing and Scheduling was held at The George Washington University in Washington, DC on February 3, 4, 1982. Its purpose was to facilitate an exchange of ideas, methods, and results concerning this subject among academic, military, and commercial researchers and analysts. An important objective of the Symposium was to assess present and future directions, models, and solution methods for this class of planning and scheduling problems. This volume contains the edited transcripts of what		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

20. Abstract (cont'd)

> was actually said at the Symposium; it presents the ideas and discussions of a group of dedicated people attempting to make progress in the resolution of a class of complex problems.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

THE GEORGE WASHINGTON UNIVERSITY
School of Engineering and Applied Science
Institute for Management Science and Engineering

Abstract
of
Serial TM-69250
December 15, 1982

PROCEEDINGS OF A SYMPOSIUM ON
CARGO SHIP ROUTING AND SCHEDULING

Washington, DC, February 3,4, 1982

Edited by
Richard M. Soland

A Symposium on Cargo Ship Routing and Scheduling was held at The George Washington University in Washington, DC on February 3,4, 1982. Its purpose was to facilitate an exchange of ideas, methods, and results concerning this subject among academic, military, and commercial researchers and analysts. An important objective of the Symposium was to assess present and future directions, models, and solution methods for this class of planning and scheduling problems. This volume contains the edited transcripts of what was actually said at the Symposium; it presents the ideas and discussions of a group of dedicated people attempting to make progress in the resolution of a class of complex problems.



Supported by
Grant N00014-82-G-0016
Office of Naval Research

Accession For	
NTIS GRA&I <input checked="" type="checkbox"/>	
DTIC TAB <input type="checkbox"/>	
Unannounced <input type="checkbox"/>	
Justification _____	
By _____	
Distribution/ _____	
Availability Codes _____	
Approved for Dist Specimen	
A	

PREFACE

A Symposium on Cargo Ship Routing and Scheduling was held at The George Washington University in Washington, DC on February 3,4, 1982. Its purpose was to facilitate an exchange of ideas, methods, and results concerning this subject among academic, military and commercial researchers and analysts. An important objective of the Symposium was to assess present and future directions, models, and solution methods for this class of planning and scheduling problems.

The contents of this volume are not proceedings in the usual sense. Rather than polished and refereed papers, this volume contains the edited transcripts of what was actually said at the Symposium; it presents the ideas and discussions of a group of dedicated people attempting to make progress in the resolution of a class of complex problems. The first day of the Symposium was devoted to five invited papers and a panel discussion. On the second day there was a sixth invited paper, and six discussion groups met, two at a time, for an hour each. These were followed by summary presentations by the discussion group leaders and then a closing panel discussion.

Financial support for the Symposium on Cargo Ship Routing and Scheduling was provided by the Maritime Administration (MARAD) of the U.S. Department of Transportation and the Military Sealift Command (MSC) of the Department of the Navy. It is a pleasure to acknowledge the organizational assistance of Walter M. Maclean and Robert O. Nevel of MARAD and Chester J. Jakowski, Jr. and Jonathan D. Kaskin of MSC. The Office of Naval Research provided contract support.

In addition to expressing my appreciation to all the participants for their enthusiastic activity in and support of the Symposium, I especially want to thank the invited speakers and the discussion group leaders. I also thank the Institute for Management Science and Engineering of The George Washington University School of Engineering and Applied Science for its support. My very special thanks for their considerable efforts go to Teresita Abacan and Dorothy Wagner for typing these proceedings and to Bettie Taggart for administrative and editorial assistance.

Richard M. Soland

Washington, DC
December 1982

TABLE OF CONTENTS

	Page No.
Department of Defense Ocean Transportation Routing and Scheduling Requirements - John H. Scott	1
Characteristics of Commercial Ocean Carrier Routing and Scheduling Problems - Russell F. Stryker	33
A Review of Cargo Ship Routing and Scheduling Models - David Ronen	48
The Utility of Heuristics in Relatively Complex Strategic Mobility Models - Carroll J. Keyfauver	82
Interactive Vessel Scheduling at EXXON - Thomas E. Baker	123
General Discussion Following the Invited Presentations of February 3	146
Discussion Group 1 - Industrial Operations - Walter E. MacLean, leader	164
Discussion Group 2 - Liner Operations - Donald K. Pollock, leader	186
Discussion Group 3 - Crisis Environment/Time Sensitive Scheduling - Joseph F. Ballou, leader	202
Discussion Group 4 - Contingency Planning and Scheduling - Lawrence A. Gallagher, leader	228
Discussion Group 5 - Exact Solution Approaches - Nicos Christofides, leader	263
Discussion Group 6 - Heuristic Solution Techniques - John J. Jarvis, leader	289
Summary Presentations of Discussion Group Leaders	314
Final Discussion Prior to the Close of the Symposium	328
References on Cargo Ship Routing and Scheduling - David Ronen, compiler	334
Symposium Program	338
List of Participants	340

Department of Defense Ocean Transportation
Routing and Scheduling Requirements

Captain John H. Scott

Military Sealift Command
Department of the Navy
Washington, DC

The Military Sealift Command (MSC) of the Department of the Navy is one of three transportation operating agencies, and it is charged with the ocean transportation for the Department of Defense. In this activity it is supported by the Military Airlift Command which, for one thing, flies our ship crews around. In coordination with the Military Traffic Management Command (for dry cargo) and the Defense Fuel Supply Center (for POL), which provide the requirements, MSC moves dry cargo and POL worldwide, during both peace and war. We also perform some non-transportation functions, which are not really relevant here.

Specifically, MSC has a five-fold mission. You see this indicated in Figure 1.1. First, it is to provide contingency and mobilization support of military forces worldwide. Note that I listed that first. There are many people who think that our first priority is the point-to-point movement of cargo in a peacetime scenario. That is not true.

Secondly, it is to provide peacetime logistic sealift in support of military forces worldwide.

Thirdly, it is to develop plans and capability for emergency expansion. It is primarily in this third mission area that this group will become interested.

Fourth, it is to operate fleet auxiliary ships for the U.S. Navy in both peace and war.

And fifth, it is to provide and operate ships for special purposes, such as scientific support. That refers to one of the non-transportation missions that I mentioned earlier.

MSC MISSION

- 0 WARTIME SUPPORT OF FORCES
- 0 PEACETIME SUPPORT OF FORCES
- 0 DEVELOP EMERGENCY PLANS AND CAPABILITY
- 0 OPERATE AUXILIARY SHIPS
- 0 PROVIDE SEALIFT FOR SPECIAL PURPOSES

FIGURE 1.1

See Figure 1.2. While MSC is both functionally and financially committed to wartime planning and execution of those plans during a crisis, it is also committed to an efficient peacetime operation. A paramount concern of the Command, however, is the smooth transition from its peacetime posture to a wartime one. As the nation's strategic sealift arm, MSC provides the quick reaction capability to deploy military forces when a conflict commences, and of course, the long-term capability to sustain those operations. We are not exactly a small business either, what with a budget of something over \$2 billion.

See Figure 1.3. At this point you might ask, what are our interests in this symposium? Put very simply, they are first, to assess recent advances in the ship scheduling and routing business. Secondly, to promote interest in model development for military purposes, and thirdly, to identify any advanced modeling which could be incorporated in my commander's decision-making process. These decisions range from determining the size of our peacetime dry cargo and tanker fleets to analyzing courses of action during crisis situations.

In order of priority, those areas of interest to MSC in model development are, first, deliberate planning. Secondly, crisis management. Thirdly, the peacetime dry cargo fleet and, fourth, the peacetime tanker fleet.

See Figure 1.4. Since the end of World War II, international crises affecting U.S. interests, which were grave enough for our government to consider using military force, have been occurring at an average rate of about every two months. The military operations in response to these crises require deployment of forces that range over a broad spectrum. At the small end of the scale, there are usually only a few thousand people involved and that would include all of the support forces, so no mobilization is necessary.

At the other end of the scale, we find our contingency plans for operations in Southeast or Southwest Asia and/or the reinforcement of Europe within the North Atlantic Treaty Organization context. The latter, of course, is the most demanding scenario contemplated for mobilization in deployment of U.S. troops. It involves over 1 million personnel and several million short tons of equipment and supplies.

FIGURE 1.2

AREAS FOR MODEL DEVELOPMENT

- o WARTIME PLANNING AND EXECUTION
- o EFFICIENT PEACETIME OPERATION

MSC SYMPOSIUM INTERESTS

- o ASSESS RECENT ADVANCES IN SHIP SCHEDULING AND
ROUTING
- o IDENTIFY MODELS APPLICABLE TO MSC
- o PROMOTE INTEREST IN MODEL DEVELOPMENT FOR MILITARY
PURPOSES.

FIGURE 1.3

BROAD SPECTRUM OF CRISIS

- o LIMITED LOGISTICAL SUPPORT OF ALLIES
- o OPERATIONAL PLANS FOR SOUTHWEST ASIA
- o REINFORCEMENT OF NATO ALLIES INVOLVING MOBILIZATION

FIGURE 1.4

See Figure 1.5. The solution to the military planning problem requires a couple of approaches. In peacetime, we conduct deliberate planning for the more lengthy and more resource-taking contingencies. In time of crisis, we determine if there is an existing plan which can be applied. If there is, we modify it and we use it. If there is no plan, then we have to perform some rather rapid time-sensitive planning.

See Figure 1.6. The deliberate planning process may be briefly described as follows: the initial phase, rightly termed "initiation," assumes a possible threat and identifies the forces available to meet that threat.

The next phase, that of concept development, determines the best course of action and develops the concept for the operation. The concept for the operation is a rather broad narrative statement of how the forces are to be allocated, deployed and then employed.

The next phase, plan development, begins to describe the requirements in more detail; in time phases and in modes of transportation. The product of this phase is a time-phased, force deployment data (TPFDD) file used to examine the feasibility of the transportation operation. The TPFDD is then turned over to the transportation operating agencies to test for feasibility in some detail. The TPFDD contains detailed feasible movement tables, a closure profile, and a delivery profile. If the plan is infeasible, adjustments are made in the priorities or in the modes of transportation until a detailed feasible movement schedule has been produced.

The next phase, that of plan review, allows the Joint Chiefs of Staff to review the plan, while the last phase, plan support, then tasks each agency which has to support the plan.

The deliberate planning process thus places a canned program on the shelf. The canned program is a feasible plan of operation to meet a defined threat.

See Figure 1.7. Now let us turn to the other approach to the military planning problem, the crisis management system. It too is composed of phases.

FIGURE 1.5

APPROACHES TO MILITARY PLANNING

- o DELIBERATE PLANNING
- o CRISIS MANAGEMENT

FIGURE 1.6

DELIBERATE PLANNING

- o INITIAL PHASE
- o CONCEPT DEVELOPMENT
- o PLAN DEVELOPMENT
- o PLAN REVIEW
- o PLAN SUPPORT

CRISIS MANAGEMENT

- o SITUATION DEVELOPMENT
- o CRISIS ASSESSMENT
- o COURSE OF ACTION DEVELOPMENT
- o DECISION PHASE
- o EXECUTION PLANNING
- o EXECUTION

FIGURE 1.7

The first phase, situation development, has as its objective the detection and assessment of events that may become a crisis.

The second phase, crisis assessment, initiates the development of the options, if the President has determined that a crisis, in fact, exists. Of course these options can be either diplomatic or military.

The third phase, the course of action development phase, is where a determination is made as to whether a suitable plan exists. We go back to the shelf and see what is available. And the support command collects the factors limiting each of the courses of action. These factors include the force data which identify the major combat forces to be employed and the total transportation assets. The TOAs prepare closure estimates for each of these courses of action.

In the decision phase, the Joint Chiefs of Staff submit the proposed courses of action to the Secretary of Defense and then, ultimately, to the President for approval.

The execution planning phase is entered when a military course of action has been decided upon. The supported commander, that is, the guy on the far shore, assisted by the deployment community, completes the force list using actual forces, origins and dates. And of course, during the execution phase, the final phase, the TOA develops the detailed movement tables and schedules.

See Figure 1.8. There are some similarities in the products required for MSC when involved in either the deliberate planning or the crisis management program. During the deliberate planning process, a gross feasibility is required by the Joint Chiefs of Staff. This corresponds to the closure estimates that are produced by MSC during the course of action development phase. The detailed movement tables produced for deliberate planning correspond to the execution planning and execution phase. Of course, as the scope of the military planning tools are described later, differences will appear, but many products are nevertheless identical.

Let us now consider the deliberate planning process in just a little bit more detail than we have thus far.

SIMILARITIES	
DELIBERATE PLANNING	CRISIS MANAGEMENT
o GROSS FEASIBILITY	o CLOSURE ESTIMATES
o DETAILED MOVEMENT TABLE	o EXECUTION PLANNING AND EXECUTION

FIGURE 1.8

See Figure 1.9. MSC is currently developing a comprehensive methodology - the short title is SEASTRAT - to evaluate the feasibility of meeting our mobilization requirements. When complete, SEASTRAT should meet MSC deliberate planning needs. The objective of the SEASTRAT model is to test the feasibility of a particular movement plan by matching ships with cargo, in order to deliver cargoes at the required ports of discharge at the prescribed times. The ship routing and scheduling must be performed consistent with the initial cargo and ship locations, port capabilities, cargo ship types, plus--and very important--efficient ship utilization designed to meet optimum delivery requirements.

See Figure 1.10. In looking at an overview of the SEASTRAT program, we note these major input categories: cargo data, load-out ports, ports of discharge, and the sort of things that you would expect to have as input to a program of this sort. The SEASTRAT model will perform ship routing, scheduling and cargo delivery, with MSC interacting with the program to specify additional constraints and special scheduling requirements that crop up during the problem.

The key output of SEASTRAT is feasibility. Are the cargo delivery requirements met? The feasibility evaluation should be available both as a gross preliminary feasibility estimate and later as a detailed ship-by-ship, cargo-by-cargo analysis of the shortfalls.

In addition, for a detailed analysis SEASTRAT should provide comprehensive cargo movement schedules, ship routing and schedules, and, very important, identification of slippage or other problem areas.

See Figure 1.11. The magnitude of the SEASTRAT program can range from 5,000 to 15,000 cargo units to be delivered during any one particular analysis. The number of available ships under MSC control could range up to 1,000 in any given scenario. There could be 150 load-out ports, and there could be up to 80 discharge ports.

The time horizon for the analysis would probably be about 180 days, which means that you could anticipate multiple voyages. Further detail of the SEASTRAT input data is shown in Figure 1.12.

The description of the cargo delivery requirements includes, again, cargo characteristics, ports of load-out and discharge, the earliest cargo

SEA STRAT OBJECTIVE

- 0 TEST FEASIBILITY OF A PARTICULAR MOVEMENT BY
MATCHING SHIPS WITH CARGO

FIGURE 1.9

SEASTRAT OVERVIEW
SHIPPING TRANSPORTATION FEASIBILITY ANALYSIS

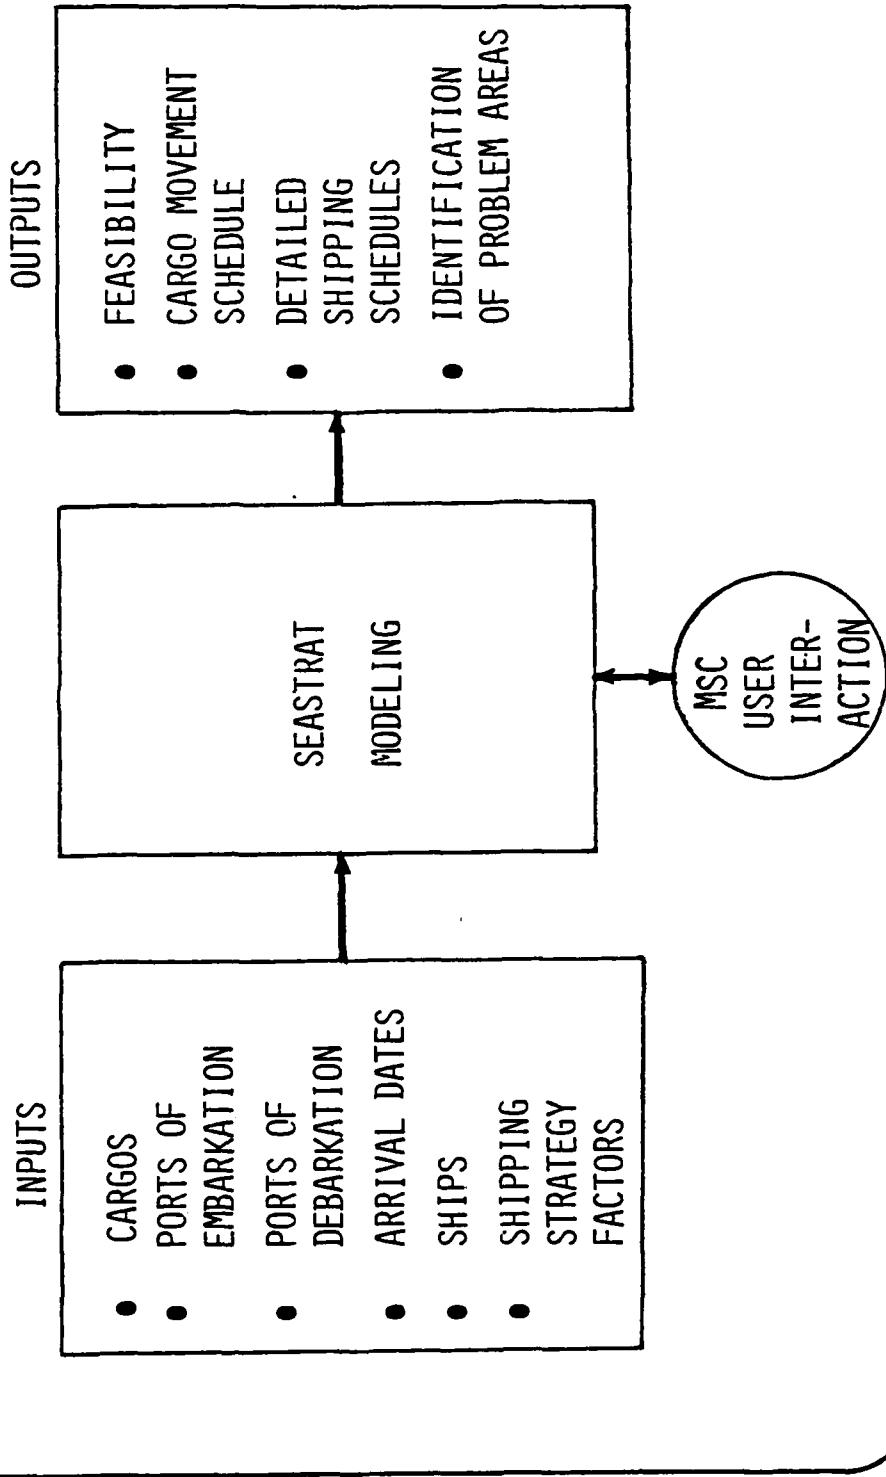


FIGURE 1.10

SEASTRAT PROBLEM SIZE

- 5,000 - 15,000 INDIVIDUAL CARGO UNITS TO BE SHIPPED
- UP TO 1,000 SHIPS
- 150 PORTS OF EMBARKATION
- 80 PORTS OF DEBARKATION
- 180 DAYS

FIGURE 1.11

SEASTRAT INPUT

- 0 DESCRIPTION OF CARGO UNITS
- 0 DESCRIPTION OF PORTS
- 0 SHIP DATA
- 0 SHIP ROUTING DATA
- 0 SHIPPING STRATEGY

FIGURE 1.12

availability dates, and of course the required delivery dates. The description of ports includes the obvious things; the ships' data base would include the obvious sorts of things. Ship routing data includes such data as recommended shipping lanes and whether or not ocean clearance has been granted for these specified lanes.

In addition there are special shipping strategy factors, such as convoying, attrition, cargo unit cohesion--that refers to the fact that certain kinds of cargo have to go together--and resupply requirements.

SEASTRAT modeling requirements are shown in Figure 1.13. SEASTRAT will determine the detailed feasibility of a plan by creating a movement table that assigns cargo to ships and that will result in the lowest landed cost, not only in terms of dollars but in terms of time. This movement table must be consistent with the available ships' characteristics, locations, cargo loading and off-load times, port throughput and shipping routes.

The model will produce both gross feasibility estimates and detailed feasibility analyses. The gross feasibility turnaround should be about 15 minutes or less, and the gross feasibility model will also be capable of sensitivity analyses for a given set of assumptions.

The detailed analysis should be available within three hours. The model should accommodate a statistical input of ship availability, such that the output of feasibility of a plan is given in terms of a probabilistic statement involving a point estimate and an interval.

To stimulate our thinking a little bit, Figure 1.14 shows several generic methodology issues for SEASTRAT development. The first issue is level of detail, which can range from gross feasibility to a detailed executable schedule. The gross feasibility answers such questions as: are there enough ships to lift a cargo in the time required, irrespective of such things as port throughput, convoying, and some of the other factors?

An executable schedule, on the other hand, is so detailed that the cargo and ship movement could be implemented in the real world.

A second issue here is the actual scheduling algorithm to be used in SEASTRAT. Heuristic methods use reasonable logic rules to schedule ships and cargoes. The problem comes in defining what is reasonable.

SEA STRAT REQUIREMENTS

- o TEST FEASIBILITY OF MOVEMENT BY PROVIDING DETAILED
MOVEMENT TABLE
- o PRODUCE BOTH GROSS AND DETAILED TABLES
- o PRODUCE PROBABILISTIC STATEMENT OF FEASIBILITY
- o CAPABILITY OF REASONABLE TURNAROUND

FIGURE 1.13

SEASTRAT METHODOLOGY ISSUES

- o LEVEL OF DETAIL
- o SCHEDULING APPROACH
- o APPROACHES FOR IMPROVING TRACTABILITY
- o SPECIAL FEATURES

FIGURE 1.14

Optimization approaches use mathematical techniques to determine the best solution; however, such techniques may not be computationally feasible for large problems. A hybrid approach could perhaps combine both heuristic and optimization methods.

An additional consideration is the need for modeling approaches which simplify the computational requirements. Aggregation can be used to combine small elements, such as an individual cargo ship's ports, into a larger element which can be treated as a single unit without altering the basic features of the problem. Decomposition techniques can be used to structure a high level large model into smaller models, which then can be solved nearly independently. For example, the individual ports could be partitioned into port regions. Approximations can be used to model complicated relationships in simpler terms, with improved approximations made subsequently as the solution is obtained.

Finally, SEASTRAT has a number of special features or needs which should be addressed in the model. These features include the need for convoying and the statistical nature of ship locations at the starting point of a mobilization scenario.

All of the foregoing issues will be addressed tomorrow in discussion group 4 on contingency planning and scheduling, moderated by Commander Gallagher.

Now let us turn to the crisis management area of the MSC mission. See Figure 1.15. The Crisis Management System will provide the support required by MSC to fulfill its operational role during a crisis or a wartime situation.

As described before, the three phases of crisis management are: first, the course of action development; second, the execution planning; and then, of course, execution. The three phases vary greatly in the quality of input.

The course of action development uses "notional" ship characteristics and availability as input and produces a feasibility estimate for a particular course of action. This compares with the execution phase which uses actual data and produces a detailed matching of ships with cargo.

FIGURE 1.15

- CRISIS MANAGEMENT
 - o COURSE OF ACTION DEVELOPMENT
 - o EXECUTION PLANNING
 - o EXECUTION

Let us quickly step through the process, under the circumstances when a crisis or contingency warrants military intervention and the Joint Chiefs of Staff (JCS) issue a crisis alert order. See Figure 1.16. MSC at this point enters the course of action phase of its Crisis Management System. The purpose of this phase is to provide a quick estimate of our sealift capability in terms of the JCS sealift requirement. The object is to evaluate the various Commanders-in-Chief (CINC) proposed courses of action, using ship availability assumptions. Ships are available in a time-phased augmentation plan which starts with the MSC-controlled fleet, and may eventually end up with the acquisition of privately owned shipping.

The model is required to produce the earliest closure date with the given resources. The model must be extremely flexible. The number of ships could range from 1 to 600, and the pool of ships could range anywhere from 40 to 1200.

The POEs and PODs could range from 1 to 150 or 1 to 80, respectively. The time horizon could range from two weeks to six months. The outputs of the model must be threefold: first, it must give a closure date; second, it must estimate the cost of the plan; and, third, the closure date has to be achievable.

And, last but not least, the decision maker wants a response within an hour. The model hopefully would be test run three or four times a year during the military exercises that we run; and, of course, the model has to consider the same port and ship constraints as the deliberate planning model did.

The next phase of the Crisis Management System is execution planning. See Figure 1.17. A selection has been made of the course of action to be followed at this point. The requirement at this stage will be less "notional" and will start selecting specific candidate ships.

The decision makers might require that the closure dates be given for various aggregates of cargo, e.g., the closure of the battalion.

The final phase is the execution of the plan. See Figure 1.18. The purpose of this phase is to manage the shipping assets of MSC during this operation, and the object of the phase is to provide detailed schedules which will match cargo with ships, such that the required delivery dates are met, and at the least cost.

COURSE OF ACTION DEVELOPMENT

- 0 EARLIEST CLOSURE DATE
- 0 COST OF SEALIFT
- 0 CLOSURE DATA MUST BE ACHIEVABLE

FIGURE 1.16

FIGURE 1.17

- EXECUTION PLANNING**
- 0 BEGIN SELECTION OF CANDIDATE SHIPS
- 0 DETAILED CLOSURE DATES

EXECUTION

- o MANAGE SHIPPING ASSESS
- o PROVIDE DETAILED SCHEDULES

FIGURE 1.18

Tomorrow's discussion group 3 on crisis management, which will be moderated by Mr. Ballou, will attempt to come to grips with the methodology necessary for building the appropriate scheduling models.

Now let us turn to one area of our peacetime operation. See Figure 1.19. The purpose of the MSC dry cargo fleet is to lift ammunition, vehicles, and other military requirements, plus household goods in support of our military forces worldwide, and to serve as a nucleus for extension upon mobilization.

The object in fleet sizing and scheduling is to meet these requirements with the optimum number of ships. By fleet sizing we mean estimating the number of ships needed to satisfy the lift requirement for any upcoming fiscal year. An objective of the fleet sizing drill is to size the fleet to meet all of the cargo requirements at least cost, while achieving the desired readiness posture. In response to the budget cycle, fleet sizing is required annually, and we usually have a couple of weeks to obtain the results required.

Also, in sizing the fleet, the need for some ships to be maintained in a reduced operating status to permit rapid response in case of an unexpected crisis has to be considered. Now, since MSC dry cargo service is normally available only where commercial service is not available, the schedule model must provide the flexibility to increase or decrease the level of service in response to unforeseen changes in commercial carriage availability.

Last year, the MSC-controlled fleet lifted 3 million tons of general dry cargo on 40 government-owned or chartered ships. Six of the 40 ships operate in a liner-type service, and the others are pretty much in a tramp mode. To meet the lift requirement, we maintain a fleet of break-bulkers that ply between about 90 ports worldwide.

The last area of our peacetime business is the tanker fleet. Again, we need some modeling development here. Again, it is in terms of fleet sizing and scheduling. MSC lifts 7 million barrels of petroleum products each month. At any instant, we have in transit roughly 3 million barrels. These movements are now made possible using about 30 tankers operating between about 140 ports.

Present MSC policy is that 95 percent of the lift requirements, i.e., the POL lift requirements, be met by government-owned or long-term chartered shipping.

MSC PEACETIME FLEET

- O SUPPORT MILITARY FORCES WORLDWIDE**
- O SERVE AS NUCLEUS FOR EXPANSION**
- O PROVIDE SERVICE WHERE COMMERCIAL SERVICE IS NOT AVAILABLE**

FIGURE 1.19

Tanker scheduling can be viewed as matching shipping assets with cargo requirements generated by the DOD community. The scope of the problem is similar to that of the cargo fleet sizing. There is the additional consideration that there are 20 types of POL products that have to be moved, and that on any one voyage two or more products can be lifted.

While the scheduler normally views lift requirements 60 days ahead, he must also schedule special missions for perhaps a year in advance. These special missions involve such things as Arctic and Antarctic resupply, wherein we use ice-strengthened tankers, or even laying out the tank ships that we propose for various military exercise programs. The scheduler then has to make daily use of the model, and it would be nice if he could have a response within two hours.

Any model must at least consider the following constraints: a tanker can only be assigned after considering such port restrictions as draft, dead-weight and length; cargo assignment can only be made after considering the last cargo lifted, due to the possibility of contamination; any model must meet all product lift requirements. The Navy has a unique tanker requirement which also must be considered, namely, the refueling of combatants and auxiliaries at sea.

That completes the description of the scheduling areas in which we at MSC feel we require some assistance in the development of models. The models for deliberate planning and crisis management are particularly challenging because of the scope of the problem and the need for the very timely response. I think what is required is the best utilization of our resources that we can possibly achieve.

And the final question that I raise is this: are optimal or near-optimal movement tables possible to meet the time-sensitive planning need for military decision making?

DISCUSSION

TEMMLER: I had the idea it was all one-way movement: is that correct?

SCOTT: Are we talking peace or war?

TEMMLER: Either.

SCOTT: In our peacetime business it is not a one-way movement. We move back to CONUS 25 percent of the tonnage that we move out. Now, in wartime, obviously, the majority of it is one-way, so you have to keep the two different scenarios in mind - whether you are at peace or at war.

WEBSTER: Do the U. S. flag carriers keep the Military Sealift Command apprised of the location and availability of their ships?

SCOTT: There is a reporting system. On the East Coast there is an organization in Norfolk called SACLANT (Supreme Allied Commander-Atlantic). They know where the merchant ships are, not only the U. S. flag carriers, but all the rest of them also.

RONEN: Who is developing the models? What state are they in now?

WILMES: I am an Army officer with the Military Sealift Command, Sealift Strategic Mobility.

We have a current model called SEACOP that was developed some 10 years ago, with 20-year-old technology. It does scheduling. At that time it was developed in order to run one operation planning exercise. It is used for deliberate planning. One has to take into account our national posture when the model was developed. We have changed our views of the world, and we are looking at more operations plans. We need a rapid response capability.

We are looking at, in my area, SEA STRAT, which is more than just a model. SEA STRAT is a management system. The word SEA STRAT does not stand for anything special; it is a name that we picked up. We are developing SEA STRAT, and we are looking forward to this group for help. We have technicians and academic people on our team, and we ask you to join our team and help us develop SEA STRAT modeling for deliberate planning.

One of the challenges we must deal with is to be able to make the transition from deliberate planning, where the analyst looks out to the future on how we would go to war, to crisis management, where the operator sits down and says, "This is day zero and we have to go to war." We need a transition to execution planning.

How far along are we? I think we have come quite a way. We have learned a lot about SEACOP. We have learned what it can and cannot do. In our particular analysis shop we have a lot of young thinkers who are suggesting things that they would like to see, and the new commander has questions about what, in turn, the Joint Chiefs of Staff want.

I have said a lot around your detailed question. Would you like to ask more about the state of the art? Is that what you are interested in, where we are in the modeling?

RONEN: At what stage of the development of the models are you?

WILMES: We are now at the point of selecting models. We have looked at some, and as far as I am concerned, I am looking to this audience to tell me, through my people and the contractors that are working for me, what is the state of the art that we can apply to the Military Sealift Command's problem, which is strategic sealift.

I guess I am tossing it back to you, saying we know where we are with SEACOP, the old model. We have looked at other models. They do not meet our requirements. We ask you to tell us what the state of the art is and what different modeling techniques we could use that are either ongoing or that could be brought online in a very short time.

KASKIN: I will say a few words about the other areas. With respect to crisis management support, we have no analytical or nonanalytical models. We produce schedules manually, based on information that we have available. It has only been in the last few years that people have really requested reasonable results, and it was determined during exercises that we really are not doing that kind of detailed scheduling. Essentially, we have not achieved any progress in developing automated tools in the crisis management area. It is still manual.

With respect to our peacetime fleet, our dry cargo fleet, we do put together a file of statistics which are used to manually carry out the fleet sizing process, but there is no modeling approach in that area. With respect to scheduling, there is absolutely no modeling, and there are no current efforts to develop such. With tankers, they do a little more detailed analysis in their fleet sizing; they try to convert the requirements on an annual basis into the handy sized tankers, and try to size a fleet with that. However,

there is really no modeling technique involved. With respect to scheduling the tankers on an operational basis, there is no modeling used. It is still done manually, and there are no developments projected in this area at this point.

Several years ago we did attempt to use an automated model, but it was not successful at that time.

WILMES: In my area, in the planning area, we are looking for ideas. We have laid out our program. We do have funds to accomplish our task. We have been approved by the Navy to accomplish the development of SEASTRAT. There is a sizable amount of money to do a task in a very short time. There are some areas in which we will need help. We are not looking for any long-term contracts.

SOLAND: I would like to make an observation, based upon Captain Scott's presentation. I think it may be something that pervades some of the other discussions and presentations today and tomorrow.

We, in mathematical modeling, and especially in the optimization area, first talk about optimization--that means an objective function. We also talk about feasibility. Considering some of the types of problems that Captain Scott talked about, I think there is a very good question as to what is the appropriate objective function or functions. Let us think about that.

I think we ought to think about what we mean by "feasibility." We talked about feasibility in terms of meeting certain schedule requirements that have been set up. But if you come close to meeting these, is that good enough? And what is close enough?

So, we ought to think some about feasibility. And these are obviously very important aspects of any kind of mathematical model, even if it is not really an optimization model. I think these are things that should be addressed during these two days.

Characteristics of Commercial Ocean Carrier
Routing and Scheduling Problems

Russell F. Stryker
Maritime Administration
U.S. Department of Transportation
Washington, DC

My purpose today is to describe briefly the several kinds of commercial ocean shipping operations and the extent to which they may be amenable to fruitful applications of the techniques of operations research, particularly optimization techniques. My primary concern is with the routing and scheduling of ships. I will not discuss OR applications to "total distribution" problems which focus on the interface between inventory analysis and transportation analysis.

In the commercial world there are two basic types of ship operations. One encompasses liner activity, which involves operation according to regular schedules on stipulated routes. It is a common carrier type of service. The other type of operation involves unscheduled service, which is sometimes referred to as tramp service. That type of service is characteristic of bulk trades, i.e., tanker trades and trades in dry bulk.

In U.S.-foreign trade, the liner trades account for about 60 million tons a year, whereas the two bulk trades account for 600 million tons a year, and that number is growing, roughly split between the liquid and the dry bulk trades.

With respect to liner trades, operation did not lend itself to optimization techniques prior to the mid-1960s. The old, common, ordinary freighter that operated until that time was a relatively slow vessel. Prior to the 1950s, speed increased to the 20 knot range, and they have never exceeded the 20-25 knot range. They characteristically spent quite a lot of time in port loading and unloading.

Breakbulk liners characteristically operate on fairly long itineraries, making several port calls on each side of the Atlantic and the Pacific. They also make a lot of port calls in both directions on the North-South routes. We still have roughly 100 vessels operating in the breakbulk mode in the U.S. flag fleet. Most of them now operate on the North-South runs.

Although analysis can be helpful in breakbulk liner ship scheduling, as, for instance, in eliminating unprofitable port calls on long itineraries, the payoff from optimization is very small in comparison with the payoff from optimizing modern, intermodal ship operations.

The shift to intermodal types began in earnest in the mid-1960s. The first and most important intermodal type is the cellular containership. The second is the barge carrier, which is similar to the cellular containership, except that it carries barges instead of boxes. The third type is the roll-on, roll-off ship, which is sometimes called a van ship and can function as a containership in its own way.

That shift provided many advantages. Perhaps the primary advantage was in terms of vessel utilization. Old breakbulk carriers were characteristically replaced by containerships somewhat larger, considerably faster, and with much better turnaround characteristics. Comparing a breakbulker with a containership on a North Atlantic route with one port call on each side, the breakbulker, at approximately 18 knots will make 13 voyages a year. The containership of equivalent tonnage, at about 25 knots, and with much faster port turnaround, will make 21 voyages a year. The containership requires one day in port versus roughly five days for the breakbulker.

Another big advantage lies in the door-to-door potential of the container. In some cases the container can be loaded at an inland origin, carried to the port, loaded aboard the vessel, and, in a very short time, carried to its foreign destination. Pilferage is minimized, and cargo handling and packaging costs drastically reduced. Everything is tightened up, and there are inventory savings of an appreciable degree as delivery times are cut.

The achievement of these benefits involves a significant cost. First of all, the ship itself is a big, expensive piece of hardware. Further, the system needs costly container port facilities with big parking areas, and crane operations are most efficient if the ship can be kept shuttling rapidly between ports, with the number of port calls minimized for the larger vessels.

To this end, container-feeder services are utilized, with smaller vessels picking up containers along the coast and bringing them to the bigger ports where the transoceanic ships are loaded and discharged. There are tradeoffs between land transportation and feeder ships in many cases. For instance, does one do better putting a container into Rotterdam and moving it by trailer across country, say, to Italy or putting it on a vessel going to one of the Italian ports? The latter would probably be slower, but the former more costly.

Here we begin to see opportunities to apply optimization techniques. This kind of situation is characteristic of most East-West routes across the Atlantic, some of those to the Mid East, and definitely those to the Far East. There are real questions of whether to run feeder service between Korea and Japan and down to Hong Kong, or whether to go directly to Korea and Hong Kong with a larger vessel. These are real questions that have had to be answered by companies operating on those routes.

The container revolution, partly because of the very high cost of the system and, partly because of its speed and sophistication, has provided an opportunity to apply many of the kinds of scheduling and optimization techniques that operations research offers to the shipping business.

The absolute benefit to be derived from the intermodal systems, and from optimizing their operation, tends to be greater in the case of short routes than long routes. If steaming time is a large fraction of total time, then the techniques that optimize operations at the port interfaces will have less to work with. Similarly, potential payoffs are greater for operators with large fleets than for those whose fleets are small.

Thus, it also offers considerable opportunity in military operations, where large fleets of freighters fall under a single operating authority. Even our largest operators, with 40-50 ships, have appreciably less opportunity for gain from the use of optimization techniques than do the Department of Defense and the Maritime Administration in operating the combined commercial fleets in wartime, when there is no concern with anti-trust considerations, and you can optimize across the entire fleet.

One of the earlier initiatives in this area was taken by the Military Traffic Management Command around 1970. They came forward then with a concept they referred to as the "dedicated port concept," under which they would dedicate a port in the continental United States to all of the cargo going to a given foreign destination, let us say New York and Bremerhaven. It has obvious advantages.

The Military Traffic Management Command, of course, operates on a least landed cost basis. Using published tariffs, a transportation officer tries to route cargo so it will arrive at its foreign destination at the lowest cost possible. However, with the cost benefits of scale that can be achieved with dedicated ports, including the benefit of very high vessel utilization, the cost factors begin to change.

The concept was offered because there was a shortage of shipping assets. However, it was greeted by significant opposition from a number of interests it would have affected, and, with the Viet Nam war winding down, it was dropped. Something like this would be worth investigating, however, for future application.

It would be impossible to apply a restrictive system such as the dedicated port system in the day-to-day world of commercial trade, characterized, as it is, by various levels of competition among ports, among ship lines, and among shipping routes. However, the liner conference system of cooperation among operators in most world trades does--where the conferences are "closed" to newcomers--provide a device through which liner capacity can be "rationalized" on given routes, that is, limited to match cargo. This permits high and efficient vessel utilization.

In the U.S.-flag liner fleet peacetime commercial ship utilization averages about 70 percent. There are various approaches to the improvement of the utilization level. One involves more effective marketing, including improvement of service. Another could involve a shift from traditional U.S. policies against closed conferences. This Administration, like preceding administrations, prefers open conferences because of the anticompetitive implications of closed conferences. However this Administration does support the reaffirmation of anti-trust immunity for participants in the U.S. trades' open conferences, plus other measures to increase conference effectiveness.

The person who finds an acceptable, legal way to increase U.S.-flag liner utilization will make a major contribution and will surely be handsomely rewarded.

Considering other potential OR applications to DOD shipping operations, nobody that I know of has looked at supply inventory savings as a function of the number of days in the surface LOC. The standard inventory models have been used to assess the value of an air line of communications as opposed to a standardized surface line of communications in terms of days of pipeline supply. However, all other things equal, surface pipelines between given points have shown relatively little variation and, hence, relatively little potential for saving. Even with the time savings afforded by containerships, there has been almost no rigorous analysis relative to the selection of surface shipping modes.

The Maritime Administration, however, has looked at the surface effect ship (SES) in this context. A surface effect ship travels somewhere between 50 and 100 knots and is capable of being built on an intercontinental size scale. It is very attractive because of its relative invulnerability to submarine attack. There has always been a small group that has promoted the acquisition by the government or the subsidization by the government of surface effect ships for that reason. We looked at potential inventory savings with respect to the SES to see whether there were high value

commodities that might be economically moved in sufficient quantity in peacetime to justify a peacetime investment in it through the government's maritime promotion program (as opposed to DOD). There are not. The amount of high-value cargo that is moved by surface displacement ships now is completely inadequate to justify the high energy cost of the surface effect ship.

The next of the intermodal ship types is the so-called barge carrier. There are between 15 and 25 in the U.S.-flag fleet. That is stated rather vaguely because some of them have been built as barge carriers and then converted into cellular containerships. There are approximately 20 right now, whereas there are 100 containerships.

The barge carrier carries big barges in cells, just as the containership carries containers. Analytically, it offers the same opportunities, but in a somewhat different way. Originally it was proposed primarily as a system to minimize port time. With the lighter-aboard-ship (LASH) type of barge carrier there is a big gantry crane that lowers a barge over the stern. The ship could leave the barge in the stream and pick it up, either empty or loaded, on its return voyage. This was expected to pay off in port cost and port time savings.

It seems not to have worked out that way. Actually the LASH system is operating most successfully where the barges move on an inland water system such as the Rhine and the Mississippi. It does not provide the same type of rapid delivery that the container system offers, and it is somewhat limited in its application.

One thing I should have mentioned with respect to containers is that the scheduling of the boxes and the barges themselves presents problems that are particularly amenable to OR solutions. There tend to be between two and three boxes or barges per cell per ship. Keeping track of these assets, and making sure they are positioned correctly, is a made-to-order OR problem, and OR techniques have been applied to it. There is probably room for additional work in this area.

The third intermodal type is the roll-on, roll-off ship. This is the ideal ship for military purposes. It is a floating garage that you can drive military equipment or outsized commercial equipment onto and off of very quickly. However, it is not very efficient commercially. It is inheriting some of the trade given up by the old fashioned freighter mentioned earlier, but we do not see a very large market for it. Maybe DOD will build a significant number of them to meet military needs, but I do not expect the industry to do that. The optimization opportunity with respect to that vessel is much the same as with the old fashioned freighter.

With respect to the non-scheduled trades, that is, the tanker trades and the dry bulk trades, there are 271 tankers in the U.S. flag fleet and 16 dry bulkers. Many of the tankers, although we call them unscheduled, actually operate on prescribed routes according to informal schedules. In a typical situation a major oil company may own or have under long-term charter as many as 80 percent of the vessels that it needs to move its crude oil and product. Those vessels will operate routinely on the same routes, according to approximately the same schedules. It is not a common user service.

In the case of the U.S.-flag fleet, there are about 15 million deadweight tons of tanker capacity, and around 11 million of that total are fully dedicated to our domestic trades, primarily to the movement of crude oil from Alaska to California and to the Gulf Coast, and to the movement of product from the Gulf Coast to the East Coast. Those ships shuttle continuously. The same ship leaves Valdez, Alaska on the 14th and, in some cases, is back at Valdez on the 25th. There is not much of an application for sophisticated analytic techniques in that type of operation.

It should also be noted that ship maintenance is much simpler than aircraft maintenance. For transport aircraft, 12 flying hours/day represents good utilization. For ships we use as a standard operating factor 350 days in the year out of 365. So scheduling maintenance, although it is a problem for the operator, does not yield a big opportunity for the application of mathematical techniques in the sense that it does for aircraft.

About 80 percent of the tonnage in the tanker industry is owned or controlled by the oil companies. The rest of the tonnage forms the so-called spot market. In that market independent operators are continuously seeking the best charters that they can get, perhaps for only a single voyage. And they try to position themselves for the best opportunity.

Theoretically, some sort of a stochastic marketing analysis might be useful to the spot market operator seeking to maximize his revenue. However, it is arguable that improved communications are more likely to help than a statistical approach to maximizing revenue in the spot market. This is an intuitive judgment--not an expert opinion.

It should be noted that all of the ocean trades are overtonnaged at this time. Perhaps the market that is best off in this regard is our own domestic tanker trade. Under the Jones Act, foreign carriers cannot operate in it. It is reserved to American operators, and the fleet is sized approximately as it should be for that trade, so the participants are not doing badly. But worldwide, where there are some 300 million deadweight tons of tankers, there have been, since 1973, a significant number of tankers in layup. Laid up tonnage reached 50 million a few years ago.

An operator facing these market conditions is in a delicate position. If he can get sufficient revenue to cover variable costs and some portion of fixed costs he will probably prefer to operate. If he cannot cover variable costs, he lays up his ship. With the present number of ships in layup, there is obviously a lot of slack in the industry. If that is the case, then what benefit can be derived from the capability to optimize scheduling?

In addition to the laid up tonnage, an equal or perhaps greater amount of tonnage is effectively slack within the operating fleet. This slack is a combination of temporary idleness which is different from layup, and of slow steaming. Slow steaming, say 12 knots for a 16 knot tanker, does two things. It provides business for additional vessels that would otherwise go into layup, and it saves money on a ship's fuel.

There are analytic approaches to the choice of correct slow speeds to minimize fuel consumption in the course of current constrained operations. Potential payoff should vary directly with fleet size.

Shifting from tankers to dry bulk carriers, there are only 16 dry bulkers in the U.S.-flag fleet. They operate in much the same mode as the tankers. The big movers of dry bulk, Du Pont, the big lumber companies, the steel companies, and so forth, have some proprietary tonnage and some tonnage under charter. Grain trade, on the other hand, accounting for a major part of U.S. bulk trade, does not involve proprietary fleets. However, major operators tend to accrue significant amounts of bulk cargo and sometimes carry it under contracts of afreightment, under which cargo is dedicated to ships.

A trade that is receiving a lot of attention now in terms of current problems and future potential is the trade in coal. Operations research techniques can be very helpful in this area. The coal trade is a growth industry. The export of coal from the United States, both to the Far East and to Europe, will probably double in the next decade. This presents an opportunity to build ships and increase employment in the maritime industry.

The Maritime Administration had done preliminary design work on a vessel referred to as an LSD, a Large Shallow Draft vessel. At about 140,000 deadweight tons this would be quite large for a dry bulker, but it could enter undredged U.S. east coast ports, which much smaller standard colliers cannot do today.

However, the LSD prospect generates a dilemma. If the approaches to east coast coal ports are dredged to accommodate standard colliers, then the cost of developing and building the LSD may not be justified. So the question becomes one of finding some way to control the balance between the construction of shallow draft vessels and the demand for dredging to accommodate existing deep draft vessels.

Another problem relates to the export of steam coal to Japan. The coal market is split between metallurgical coal and steam coal. Most

metallurgical coal is shipped from the east coast and will continue to be shipped from the east coast. Steam coal can be shipped from either the east coast or the west coast.

Between now and 1990 west coast port facilities will be unable to handle the steam coal volume expected to be exported to Japan. Therefore we expect most of that coal to move from east coast ports to Japan via the Panama Canal between now and about 1990. The large shallow draft bulker that I referred to is too wide to transit the Canal. Thus there is the prospect of a growing trade moving in relatively inefficient ships through the Panama Canal until 1990, when the loading capacity of the west coast ports will have increased to the point where export volumes will justify the use of larger, more economical vessels.

These two coal export problems appear ready made for operations research solutions--at least with respect to the phasing of shifts from low to high efficiency systems.

In summary, there are many large and small problems in the commercial ship operations business that would at least bear investigation with respect to operations research applications. However, I do not believe that the payoff from the application of optimization techniques to peacetime commercial ship routing and scheduling would be appreciable--at least in comparison with the potential payoff from their application to wartime deployment problems.

This conclusion is based in part on the relatively small sizes of the individual operators' fleets and in part on the relatively straightforward nature of the commercial scheduling problems, which lend themselves to efficient intuitive solution. In the wartime deployment case, on the other hand, hundreds of vessels are under central operational control, and scheduling/routing problems can be highly complex. Furthermore, some peacetime institutional constraints to the implementation of optimum solutions are not present in wartime.

DISCUSSION

WEBSTER: I have a comment about the Pacific container trade. I think it is important to note that Sealand and American President Lines have a head-to-head competition, but they handle their operations rather differently. Sealand handles their operation very much as you describe, with large, fast ships going between the U.S. and Japan for the feeder service. APL is an itinerant and carries their cargo to all of the individual ports, perhaps 10 or 12 ports on one itinerary. Both compete rather well.

STRYKER: Maybe they complement each other.

WEBSTER: There are plusses and minuses on both sides. The feeder line service takes longer to deliver cargo because you have to transfer the cargo, and if you are carrying a lot of California fruits and vegetables and things of that sort, that makes a difference. Carrying TV sets back from Hong Kong, which is a high value cargo, the inventory cargo time makes a difference. So what it costs you extra in running you make up in these other areas, it appears.

KASKIN: You said that in the area of the breakbulk service, other than scheduled liner service, operations research techniques might be applicable in determining the sequence of the ports and the speed of operation of the ports to take in factors of market share that can be achieved under those various sequences and speeds. Do you take into account operating costs versus those revenues and apply those techniques to look at the problem?

STRYKER: I did say that, but I meant to imply that that problem is basically trivial. A competent operator will know his costs to make a given port call. He will know steaming fuel costs, cargo handling charges, etc., and all of the ship operating costs he will incur during the period of the port call. He can readily contrast these costs with the revenue he is likely to get from the cargo that is offered at that port. A good

operator will do this in his head after some experience. Analytically, it is not a difficult problem. Modeling techniques may be applicable, but I really do not see it as an opportunity for optimization.

TEMMLER: You said all trades are overtonnaged at this time. Could you define "this time," how long a period, when it began, how long it is going to last?

STRYKER: The maritime trades are referred to as cyclical, i.e., boom and bust, and the cycles are simplistically thought of as covering decades. However, I do not necessarily subscribe to that. Let us talk about the tanker trades. The overtonnaging in the tanker trades became most apparent after the embargo of 1973. It really existed before that. Before 1973 the world oil trades were expanding rapidly. There was some overbuilding in anticipation of that market. Thus, even if the bottom had not fallen out of the market after the 1973/1974 embargo, there would have been some overtonnaging. However, at least theoretically, there would have been an opportunity, eventually, for demand to catch up with the supply of ships.

In a totally unregulated and uncoordinated industry, though, there is a real question as to whether demand ever really catches up. For instance, the conventional wisdom says that although the world bulk trades are now overtonnaged, new vessels will be needed within a few years because the trade is growing fast. I am not sure that is true.

Possibly there is a "bow wave" of construction with the entrepreneurs always looking two years ahead. The new construction always provides two tons of capacity for each ton of future cargo because the entrepreneurs are all planning to move the same tons. Thus I cannot accept conventional wisdom as it would apply to the bulk trades, which are dominant. I personally do not really foresee a situation where there is not at least a little more tonnage than is really needed. I think this generally applies also to the liner trades. The liner trades worldwide,

except for the United States, turned to the closed conference system to limit liner tonnage starting in 1880, and they institutionalized it in the period of World War I. However, the conference system is not very tight. I think there will always be a little more tonnage than is needed.

TEMMLER: The industry will have more capacity than its demand, but is that different from any other industry's method of operating?

STRYKER: It can be. Right now the U.S. industrial plant is operating at 70 percent--the steel industry is at about that level--but we have had times when it was well over 90 percent. In shipping I do not foresee it being well over 90 percent except occasionally for the liner trades, and even there it will take some changes in the U.S. law before that degree of rationalization is achieved.

BENTLEY: At a recent MARAD conference, a paper was presented describing the development of a systems dynamics model to predict the market response to changes in scheduled liner rates. To validate the model, the authors studied observed behavior on one particular route since 1940.

One important parameter in the model was the "accelerator" which quantifies the entrepreneurial response to the change in capacity in reaction to price changes in the market. Although the principal purpose of the historical study was to demonstrate that the structure of the model was reasonable, I considered the actual fitted values of the parameters to be of interest as well. The results indicated an "explosive" response to price changes, such that market behavior on the part of shipping companies, both in terms of price and capacity, would never converge to equilibrium but would tend to become ever more volatile. This research provides an empirical indication that the "bow wave" phenomenon, to which you referred, may have troughs as well as crests and, left to itself, will not settle down.

On another note, the problems of ship scheduling in tramp service must reflect somewhat different considerations than in scheduled service. Furthermore, the increasing prevalence of contract of affreightment as opposed to spot chartering must have a bearing. Would this not alter the nature of the scheduling problem?

STRYKER: If I did not say that, I intended to say it. You get to know better where the cargo is and what you are going to be able to do with it,

where it will be and what revenue it is going to generate for you. You then have the basis for the application of scheduling techniques.

BENTLEY: It makes the carrier responsible for the total distribution problem.

STRYKER: Yes.

MENTZ: One statistical fact about the dry bulk fleet which I find somewhat intriguing is that there are about 4200 dry bulk vessels currently in the world fleet. There are 800 new vessels on order. Even though it is clear that everyone understands that the current supply-demand relationship is out of balance, I find it interesting that there are 800 vessels on order at a time of substantial overtonnaging. It seems that because of these additions to the world fleet, perhaps for reasons of improved fuel efficiency in vessels of the future, perhaps because of the larger size vessels or for some different configurations of vessels like the wide-beam vessels we talked about before, there will always be large additions in capacity. There does seem to be a problem of projecting when that supply-demand balance might ever come about. And from what Russ has mentioned and the statistics, it seem that that point keeps getting further and further away.

That is a comment on the bulk side. On the liner side, besides the project that was being reported on by Mr. Bentley, which was in the marketing area, we have done some work using the system dynamics tools developed in the MIT-Cambridge Associates environment to look at the liner trade in the Pacific. We have looked at the cycles of supply and demand and certainly one of the conclusions of the system dynamics techniques is that the cycles are alway there and always will be there. The best thing you can do is accept that and try to accommodate them, and then find out how you should lead in response to those cycles as best you can, both for an individual company and for a national response. But the cycles are there and you have to recognize them.

SOLAND: I think it might be worth noting that with respect to the applicability of optimization and other OR models, anyone who reads Interfaces can attest to the fact that there are plenty of instances where substantial savings have been found by using OR and optimization models. It has been well documented. There are also cases, and Gene Woolsey likes to talk about them, where OR optimization models have fallen flat on their faces because they could hardly do any better, and sometimes worse, than what people were doing through experience and by the seats of their pants. It does not mean that we should stop trying to improve with quantitative models.

It will be interesting to hear what some of the commercial people say they perceive as problems where they think OR models can help, and whether they agree with Russ Stryker on the areas where it does not seem likely that improvement is to be found through the use of quantitative models.

A Review of Cargo Ship Routing
and Scheduling Models

David Ronen

School of Business Administration
University of Missouri-St. Louis
St. Louis, MO

We all know the importance of ocean transportation of cargoes. Many people with whom I discussed this earlier wonder why so little work has been done in cargo ship routing and scheduling. Figure 3.1 shows several explanations. First, there is the low visibility. Everybody sees trucks on the road, but in the United States it is difficult to see a ship, even in some ports. Low visibility is one major reason for the relatively small amount of work.

Second, there is a large variety in problem structure and in the decision environment in shipping, much larger than in the types of vehicle scheduling problems which have been discussed in the literature.

Third, there is more uncertainty in shipping operations, and uncertainty is not necessarily good for modeling, as everybody knows, because most modeling approaches for scheduling are deterministic and so do not address uncertainty. Actually, three Russian authors did an analysis of scheduling and found out that the probability is only about 0.3 for a ship to meet its quarterly schedule. That is not too good. One of the reasons is that ships cost a lot of money. Thus you usually do not build any slack into the operation of a ship, and once something goes wrong, all of the schedule is messed up and the ship will not meet it. I shall mention later another aspect of the uncertainty when I compare ship scheduling to vehicle scheduling.

RELATIVELY LITTLE WORK DONE DUE TO:

1. LOW VISIBILITY.
2. LARGE VARIETY IN PROBLEM STRUCTURE AND ENVIRONMENT.
3. LARGE UNCERTAINTY IN OPERATIONS.
4. VOLATILE INTERNATIONAL MARKET - STRESSES CAPITAL INVESTMENTS.
5. CONSERVATIVE INDUSTRY.

FIGURE 3,1

Then we have a volatile international market in shipping which stresses the capital investment. There is a large effect of capital investment decisions when you operate ships; the decision of what size ship to buy, what type ship to buy and when to buy it is much more important than the operational aspects where you may lose a day or two in the scheduling of the ship. You can make millions of dollars by buying and selling the right ship as a commodity and you can save only hundreds of thousands of dollars by proper scheduling. So there is much more stress on the capital investment aspects.

You know that ship owners take advantage of international laws or lack of laws. They take advantage of different taxing regulations and different crew size regulations to operate a ship in the international market. All of this has a much larger effect on the bottom line than scheduling, and therefore much more attention is directed to those areas.

In addition, the ocean shipping industry is a relatively conservative industry. The ship is probably the oldest mode of transportation except for legs and animals. It has been around over 2000 years. It is mentioned in the Bible. So there is a lot of tradition in shipping, whereas trucks have been around for only about 70 or 80 years and airplanes even less. Thus, there is a lot of resistance to change in shipping.

For example, in the airline industry there is a quantitative group which transfers information between the airlines (AGIFORS). But in many shipping companies there are hardly any quantitative analysis groups.

Before I go on let us specify the terms shown in Figure 3.2. The first term is "shipping." That usually means shipping of cargoes by ship, and not the wider context of shipping cargoes in general. "Routing" has an interesting interpretation in shipping. Very often routing means weather routing, or choosing the route between two ports which is the most efficient in terms of avoiding bad weather. In our case, routing would mean specifying a sequence of ports of call for ships.

TERMINOLOGY:

- * SHIPPING
- * ROUTING (WEATHER ROUTING)
- * SCHEDULING
- * SHORT TERM
- * MEDIUM TERM
- * LONG TERM

FIGURE 3.2

"Scheduling" is basically routing with times and dates associated with loading and unloading operations. "Short term" I would specify in cargo shipping as dealing with decisions for a very short time ahead, say several weeks, and usually in such time frame the size of the fleet cannot be changed. It is very easy to charter a ship or sublet ships, but in most cases it is not possible within 10 or 15 days. It depends very much on the market and in the Persian Gulf you can charter a ship and have it delivered within five or 10 days. But in most places it takes at least two weeks to get a ship. Short term thus means up to two or three weeks.

"Medium term" usually means up to several months ahead, and in that time range you can change the size of your fleet. You can charter and you can sublet ships.

"Long term" is whatever is beyond about half a year. That is a long time in shipping, mainly due to the uncertainty.

WEBSTER: I have a question. Generally speaking, if you are in the market to charter, so is everybody else, if it is a competitive market. Or if you are in the market to sell, so is everybody else. And it is rather unstable.

RONEN: Not necessarily. Usually corporations such as oil companies size their fleets somewhat below their long-term requirements. Eighty percent of their requirements they will fill with owned vessels and long term charters or time charters. They add capacity when they need it on the spot market, single charters for a single voyage. Actually, most of the independent owners make their bread and butter in the spot charter market.

In addition, I would like to make a definition of modes of operation in shipping. See Figure 3.3. First, there are liners, which operate similar to a bus line. There is a published timetable and an itinerary, and the ship follows the itinerary and tries to follow the timetable too. There are several differences from a bus line. One is that the ship is not necessarily empty ever; thus, you cannot always identify the beginning of the line and the end of the line. The ship carries cargoes between various ports, and it may never be empty. Some cargo may be left for the

MODES OF OPERATION:

- * LINER
- * TRAMP
- * INDUSTRIAL

FIGURE 3.3

next course. That is the major difference between the bus line and the liner.

A second type of operation is a tramp operation, which can be compared to a taxicab company. You send the ship to load the available cargo, or to where you expect cargo to be available. You basically contract for carrying a certain cargo from a loading port to an unloading port.

The third type is industrial operations. Russell Stryker lumped tramp and industrial together. I would like to differentiate between the two, because the respective operators usually have different objectives and different constraints. Industrial operations can be compared to the operation of a private fleet of trucks. The operator controls the cargo and the ships. Usually, industrial operators are found in bulk trades like oil, ore, and coal. All of the major oil companies are industrial operators. The initial reason that an industrial operator enters shipping is to give good service to his own cargo. He does not care about anything else. Eventually someone will get up and say, "Why are we running 50 percent of the mileage empty? Why do we not take outside cargo?" And in that case the operator complicates his life, and he begins to share the problems of a tramp operator. And very often he regrets it.

As I said earlier, ships can be easily chartered and sublet. That is done in an international market. There is an exchange in London where you can charter ships and sublet ships whenever you want, just like a commodity market.

Since many here are well acquainted with the vehicle scheduling problem, I would like to discuss briefly major differences between vehicle scheduling and ship scheduling problems. See Figure 3.4.

Ships are different from each other. Every ship basically has different characteristics. They have different sizes, different speeds, and if you find two ships with the same physical characteristics, they may have a different cost structure. You may have chartered them at different times, and the charter market is quite volatile, depending on the supply and demand at the specific time. Since ships are different

4

DIFFERENCES FROM VEHICLE SCHEDULING:

1. SHIPS HAVE DIFFERENT OPERATIONAL CHARACTERISTICS (SIZE, SPEED) AND COST STRUCTURES.
2. THE SCHEDULING ENVIRONMENTS DEPEND ON THE MODE OF OPERATION.
3. SHIPS DO NOT NECESSARILY RETURN TO THEIR ORIGIN.
4. HIGHER UNCERTAINTY IN OPERATIONS.

FIGURE 3.4

in size, treatment of supply and demand larger than the shipload is not simple. You do not know how many shiploads you have until you specify the ships.

The scheduling environment of ships depends very much on the mode of operation. Vehicle scheduling problems resemble industrial operations, but tramp and liner operations are quite different. Ships do not necessarily return to their origin. That is a very common assumption in vehicle scheduling. Ships which you take on charter may have to be delivered to certain geographical zones. Ships go once or twice a year to a yard for maintenance. Basically, you cannot count on the fact that ships will return to the origin.

There is higher uncertainty in shipping operations. As I mentioned earlier, there is uncertainty that stems from many sources. We have the ship itself, which is a mechanical system. We have weather conditions. We have loading and unloading problems. The quality of manpower on ships and on the shore is not the best. You see many strikes in those places, and slow downs too.

There is an additional difference. When we deal with vehicle scheduling, we speak usually about a daily schedule, and the night is used as a buffer. If a truck is delayed during the day, it will work two hours more, but at night the schedule is finished, and there is no delay in beginning the next day. The night is used as a buffer against uncertainty, against accumulation of workload. Ships work 24 hours around the clock, and there is no buffer against uncertainty, so all delays are accumulated.

EMBERGER: Not with over-the-road trucks. You are talking about local delivery type truck operations within a very small geographical location. You are not talking about a moving company truck, where the driver may never return to the place of origin.

RONEN: That is definitely true. But most vehicle scheduling models deal with the local delivery problem.

EMBERGER: It seems like you are comparing apples and oranges.

WEBSTER: Actually, cargo liner trades usually have a little bit of buffer space, because you wish to arrive in most ports say, on week days rather than weekends. You do not want to arrive at night. You want to arrive in the morning, and so forth. There are several legs of the route which you travel at a lower speed, purposely. That does provide the buffer.

RONEN: You also want to leave the port before the weekend; right?

WEBSTER: There is a buffer built in, because you lower your speed in order to achieve that.

RONEN: These are some mechanisms to address buffers. I do not know to what extent they really help scheduling.

DOUGLAS: I think you exaggerated perhaps on the higher uncertainty in operations, when you referred to the difference in quality of manpower. You referred to differences in quality between shore operations, sea operations, and truck operations. I have a little problem with that.

RONEN: That is my feeling; I do not have data to support it. Another difference, and this may be a minor difference, is that a ship may be diverted to a different unloading port when it is loaded and sailing at sea, and I do not know to what extent this happens in truck operations. This is true especially for bulk ships and major bulk commodities, where the commodity is usually a uniform commodity.

I am going to skip the next section, which deals with a classification of problems and models, and is shown in Figure 3.5.

Let me discuss different types of models. I have divided them into four basic types.

See Figure 3.6. First, there are transportation system models, which fall under long-term planning. In these models the sea shipping leg is only a component of the total system. Basically, the design of a transportation system produces certain constraints on the ship scheduling and the routing environment. It may be the fleet size, the fleet mix, the origin ports, the destination ports, or the storage capacities in the ports. All of these are constraints on ship scheduling

6

CLASSIFICATION OF PROBLEMS AND MODELS:

- A. MODE OF OPERATION
 - 1. LINER
 - 2. TRAMP
 - 3. INDUSTRIAL
- B. LOADING AND DISCHARGE TIMES
 - 1. SPECIFIED (SHIP SCHEDULING PROBLEM)
 - 2. TIME WINDOWS
 - 3. OPEN (ROUTING PROBLEM)
- C. NUMBER OF ORIGINS
 - 1. ONE
 - 2. MULTIPLE
- D. NUMBER OF DISCHARGING POINTS
 - 1. ONE
 - 2. MULTIPLE
- E. NUMBER OF LOADING PORTS PER VESSEL VOYAGE
 - 1. ONE
 - 2. MULTIPLE
- F. NUMBER OF DISCHARGING PORTS PER VESSEL VOYAGE
 - 1. ONE
 - 2. MULTIPLE

FIGURE 3.5

- G. NUMBER OF COMMODITIES
 - 1. ONE
 - 2. MULTIPLE
- H. FLEET SIZE
 - 1. ONE VESSEL
 - 2. MULTIPLE VESSELS
- I. TYPES OF VESSELS
 - 1. ONE
 - 2. MULTIPLE
- J. DEMANDS (SHIPMENT SIZES)
 - 1. DETERMINISTIC
 - 2. STOCHASTIC
 - A. CONTINUOUS
 - B. DISCRETE
 - 3. DECISION VARIABLES
- K. CRUISING SPEED AS A DECISION VARIABLE
 - 1. YES
 - 2. NO

FIGURE 3.5 (CONTINUED)

L. FLEET SIZE AND COMPOSITION

1. SPECIFIED - CANNOT BE CHANGED (SHORT TERM PROBLEM)
 - A. SUFFICIENT
 - B. DROP CARGOES
2. CAN BE CHANGED (MEDIUM TERM PROBLEM)
3. CONSTANT OVER SCHEDULING PERIOD
4. CHANGES OVER SCHEDULING PERIOD

M. PORT ENTRY CONSTRAINTS ON VESSELS

1. EXIST
2. NONE

N. SEA ROUTE CONSTRAINTS ON VESSELS

1. EXIST
2. NONE

O. PORTS PRECEDENCE REQUIREMENTS

1. EXIST
2. NONE

FIGURE 3.5 (CONTINUED)

P. COSTS

1. FIXED COSTS (OPERATING AND CAPITAL)
 - A. IN OPERATION
 - B. IN LAYUP
 - C. CHANGE OF STATUS COST
2. VARIABLE COSTS
 - A. STEAMING COSTS
 - B. PORT ENTRY CHARGES
 - C. TIME IN PORTS
 - D. UNIT SHIPPING COST
 - E. DEMURRAGE

Q. OBJECTIVE

1. MINIMIZE COSTS
2. MAXIMIZE PROFITS
3. MAXIMIZE UTILITY

R. CARGO TRANSSHIPMENT

1. ALLOWED
2. EXCLUDED

S. TIMES BETWEEN EVENTS

1. DETERMINISTIC
2. STOCHASTIC

T. OTHER PROBLEM SPECIFIC CHARACTERISTICS

FIGURE 3.5 (CONTINUED)

TRANSPORTATION SYSTEMS MODELS
(LONG TERM PLANNING)

PARTS OF THE PURPOSE OF THE MODELS IS TO SELECT THE PROPER TYPES
OF SHIPS (FLEET SIZE AND MIX) AND TO ASSIGN THE SHIPS TO ROUTES, -
MATH. PROG.

- * CONLEY, FARNSWORTH, KOENIGSBERG AND WIERSMA (1968)
- * NASLUND (1970)
- * MATHIS (1972)

FIGURE 3.6

and routing. There is actually some hierarchical decision making, and the question is to what extent these constraints can be taken as such or should be part of the whole system design problem.

Several models have tried to combine routing of ships with system design. Usually, only industrial operators can do this, can incorporate system design with routing decisions, because they control both the transportation system and the routing of the ships. Liner operators do not have such a control. They can decide what ports to call at, but they do not design the transportation system. The same is true with tramp operators.

The approach to these problems is usually via mathematical programming models. The first model was by Conley, Farnsworth, Koenigsberg and Wiersema in 1968. They dealt with the supply of homogeneous products from overseas to 400 inland U.S. destinations with a fleet of 50 ships of six types. They had to decide what ports to unload at and what ports would supply the inland destinations, i.e., the assigning of inland destinations to unloading ports. They used a mathematical programming model, and they wanted to minimize their costs. Part of the decision was how to assign the ships to the routes, and what types of ships would serve each route.

They simplified the problem by disconnecting the link between the land transportation segment and the ocean portion. They specified a set of U.S. unloading ports, and assigned ships to routes with the given set of ports. They did this for many different sets of unloading ports. For each set of unloading ports they dealt with the land transportation portion, the assignment of inland destinations to the unloading ports.

The next model was discussed by Naslund. This was a distribution system in Northern Europe. There were several loading ports which were specified, where wood pulp was available. There were many alternatives for unloading ports in Northern Europe, mainly in England. The objective was to minimize transportation costs from the ports of origin to the inland destinations. He approximated the cost functions in a manner which allowed him to use a procedure similar to the Baumol and Wolfe

procedure for locating warehouses. That procedure is a heuristic embedded in mathematical programming. It helped to choose the unloading ports, and gave a local optimum but not a global one. He used analytical continuous cost functions, and that raises a question of the validity of these cost functions and the heuristic results.

The last model deals with a whole distribution system. It was built by Mathis. He dealt with the problem of an oil company which has certain sources of oil and certain destinations. He made some restrictive assumptions: only one type of crude was allowed, tankers were assigned to specific routes (they did not change routes). Moreover, the sizes of the tankers were continuous. That means you can build any size tanker. Here again, if you assume a continuous variable for the size of a tanker, you have continuous cost functions for building tankers and probably to operate them. He wanted to minimize the cost of the system. He included also the cost of storage facilities on shore and the size of the tanks. And he allowed transshipments of cargo only in prespecified unloading ports.

The problem was formulated as a nonlinear integer program. He took advantage of the structure of the problem and solved it by a branch-and-bound procedure. This was a dissertation done at Case Western Reserve University. He looked at what oil companies were doing to solve the problem; they usually used simulation models.

The next type of model is for liner operations, and you do not find very many models. See Figure 3.7. First, the demand for liners depends very much on the service provided. That means the demand depends on the frequency of the service, the speed of the service, the reliability of the service, and therefore the scheduling decisions affect the demand for the service. It may be hard to quantify or put in mathematical form, but there is definitely a relationship between the service, the schedule and the revenues generated.

The objective of a liner operation will usually be to maximize profit per time unit. Again, there is uncertainty, both in the liner operations and from demand. There is some uncertainty about the demand

LINER OPERATIONS

- * DEMAND DEPENDENT ON SERVICE
- * MAX. PROFIT/TIME
- * UNCERTAINTY FROM OPERATIONS AND DEMAND
- * SIMULATION

ALMOGY & LEVIN (1970) - CARGO SELECTION

NEMHAUSER & YU (1972) - FREQUENCY AND DEMAND
RELATIONSHIP (RAIL)

BOFFEY, EDMOND, HINXMAN AND PURSGLOVE (1979) - CONTAINER
LINE SCHEDULING, SIMULATION & HEURISTICS

FIGURE 3.7

for the services which increases the difficulty of the mathematical modeling of liner operations. Most of the models of liner operations are simulation models.

The problem is complex, both due to the dependence of demand on service parameters and due to the uncertainty. Simulation has been a major mode of treating liner operations. Olson, Sorenson and Sullivan used a deterministic simulation model to evaluate scheduling alternatives in a liner company operating between the U.S. West Coast and Hawaii.

Kyland used a more elaborate model. It was a stochastic simulation model with linear programming to simulate liner operations.

POLLOCK: You made a statement that the operating objective is to maximize profit per unit of time. Would you explain that a little, what you mean by maximizing profit per unit of time?

RONEN: On one hand there are revenues, and on the other hand there are costs. You would like to increase the difference between the revenues and the costs as much as possible. It is not the same having \$1 million profit over 10 years as over six months. Therefore, you would like to increase as much as possible the difference between revenues and costs on the basis of time. You can increase revenues to some extent by increasing the frequency of calls in certain ports or adding other ports, but on the other hand this will increase costs. There is a relationship between revenues and costs. Revenues and costs are not independent of each other, and you would like to maximize the difference between them, per time unit.

WEBSTER: It would be more precise to say that the liner operator is trying to maximize his return on investment. I think it is a little bit different than to maximize profit.

RONEN: If the investment is constant, maximizing return on investment is the same as maximizing profit per time unit.

VOICE: There are some possibilities that may be coming up in the laws, like the Slade Gorton bill, that will permit the possibility of doing

something called revenue sharing or cargo pooling, which may make it possible to look at some other alternatives with respect to the way that you are looking at the problems.

RONEN: Yes. You are talking about conferences which exist in many trade routes around the world.

VOICE: That is about to become law in the United States. It is proposed.

RONEN: Several models look at specific aspects of liner operations. One is by Almogy and Levin. They took a more rigorous approach to a narrower problem. The question is what cargo to select out of the cargoes available. This is a stochastic model with a separable objective function. Each cargo has its own unloading port, and this affects the routing. Their model is quite complex. They did not present any examples.

The next model was developed by Nemhauser and Yu.

BALLOU: The optimization is the maximization of profit per unit of time?

RONEN: Yes. Nemhauser and Yu developed a model for line service. It is for commuter trains in urban transportation. They related the frequency of service to the demand. The demand and the frequency of service are dependent, and the timing of the service, is a related factor too. They use dynamic programming to find an optimal schedule, which maximizes profit over the planning horizon. Although their model was developed for train service, it can be applied to liner shipping.

The latest work in liner shipping is by Boffey, Edmond, Hinxman and Pursglove. They dealt with a container line service across the North Atlantic, between Northern Europe and the U.S. East Coast.

What they built is more than a simulation model, an interactive model which gives the user two options. The user can either specify a schedule and be able to see the financial results of the schedule in front of him on the screen, or he can use certain heuristics built into the model to generate a schedule, and then either use the schedule or modify it manually. They also gave a very good short discussion about the realities of scheduling. For scheduling they used a heuristic that looks only one step forward, a greedy heuristic.

The next area is tramp shipping (see Figure 3.8). Tramp shipping involves ships that are controlled by their owners, and they are looking for cargo, like a taxicab operation. In tramp shipping, there are many owners, and most of them are small ones which own several ships. Very few owners are big ones.

The segmentation of that part of the industry has resulted in the fact that hardly any work has been done in this area. The only work I could come up with was done by Appelgren, a Swede. He worked for a large consortium of refrigerated tramp shippers. The trade in refrigerated cargo is interesting. It is highly seasonal and it involves large quantities. No operator will buy refrigerated ships to operate them for four months out of the year. Therefore, carrying refrigerated cargoes is one of the major operations in tramp shipping. The problem treated was as follows. The operator has a given fleet, and he has a given set of cargoes that must be loaded during the planning horizon.

There are optional cargoes that may appear during the planning period. The planning period was about four to six months. The operator can either take optional cargoes or refuse them.

A cargo is a full shipload which may have one, two, or maybe three unloading ports. The objective is to maximize the revenues minus the costs, i.e., profits of the optional cargoes.

This model used mathematical programming. Actually, there were two articles. The first one used mathematical programming; it was a 0-1 linear programming model, and the Dantzig-Wolfe decomposition was used to solve the problem. The problems turned out to be network flow problems. But integer solutions were not assured, and they were rounded off. The second paper addressed the problem of rounding off non-integer solutions. It used a branch-and-bound method to address that problem.

The fourth type of model is for industrial operations (see Figure 3.9). Usually such models deal with bulk or semi-bulk commodities, such as oil, ore, coal, grain, lumber, pulp, fruit, sugar, and potash.

TRAMP SHIPPING

- * MANY SMALL OPERATORS
- * APPELGRAN (1971, 1969):
 - GIVEN FLEET
 - GIVEN SET OF CARGOES THAT MUST BE LOADED,
(LOADING TIME WINDOW)
 - OPTIONAL CARGOES MAY BE REJECTED
 - ONE CARGO PER TRIP
 - MAX. (REVENUE - COST) OF OPTIONAL CARGOES
 - MATH. PROG.

FIGURE 3.8

INDUSTRIAL OPERATIONS

* BULK AND SEMI-BULK COMMODITIES

TANKER SCHEDULING (MATH. PROG.):

DANTZIG AND FULKERSON (1954) - IDENTICAL TANKERS, LOAD AND DISCHARGE DATES, ONE LOAD AND ONE DISCHARGE PORT PER VOYAGE.
MIN. # TANKERS.

FLOOD (1954) - GIVEN FLEET, MIN. BALLAST DISTANCE.

BRISKIN (1966) - SEVERAL DISCHARGE PORTS.

BELLMORE, BENNINGTON AND LUBORE (1971) - DIFFERENT TYPES OF TANKERS, PARTIALLY LOADED TANKERS, TIME WINDOWS, MAX. UTILITY.

MC KAY AND HARTLEY (1974) - MULTIPLE PRODUCTS.

FIGURE 3.9

The military establishment in this country faced a tanker scheduling problem, and it was originally treated by Dantzig and Fulkerson. It was after the Second World War, during which a fleet of 16,000-ton tankers was built (T2), and there were hundreds of them. So, at that time the Navy had a fleet of identical tankers. They had load and discharge dates. Each voyage had one loading port and one discharging port. Dantzig and Fulkerson tried to minimize the number of tankers needed to meet those loading and discharge dates. They solved it by using a transportation problem, and deriving the schedule from the results.

Flood, around the same time, looked at a different aspect of the problem. Dantzig and Fulkerson tried to minimize the size of the fleet, but Flood took a given fleet and tried to minimize the distance ships steam in ballast, i.e., minimize the empty distance traveled. He too used a transportation model.

Briskin expanded the model a little. He allowed several discharging ports instead of one. Actually, he went back to the former context by aggregating discharging ports in a certain zone to one port and scheduling the ships and then, by dynamic programming, routing the ships through their discharging ports.

Bellmore, Bennington, and Lubore took a much wider view of this problem and added more realistic aspects by allowing different types of tankers. They allowed partially loaded tankers. Instead of having specified the load and discharge date, they allowed time windows so that loading and discharge could be done within a time window. And they tried to maximize utility. Now, the question is: What is "utility?" They had an elaborate definition of utility. There was a negative utility assigned to having more ships introduced into service and also to empty legs of trips. There was positive utility assigned to making deliveries. And, it was not assured that all of the deliveries would be made. They proposed a solution procedure that used decomposition with a branch-and-bound algorithm, but they did not show how to apply their procedure. They did not show any results. And certain people claim that their proposed approach has never been applied.

The last and most comprehensive treatment of this tanker scheduling problem was done by McKay and Hartley. In addition to all of the aspects mentioned above they allowed multiple products. Moreover, they were trying to minimize the cost of buying the oil and shipping it. There were options to buy the oil at different ports. So, they were looking at the total cost of buying and shipping oil, and they used a linear program with a specialized rounding procedure for the treatment of integer variables.

Again, there was no optimal solution. Their model seems most realistic. But, as far as I understand, it has not been applied. This is the present status of the tanker scheduling problem, but we have quite a few other industrial operations models, as shown in Figure 3.10.

One, treated by Laderman, Gleiberman, and Egan, was a routing problem. No times were used for loading and unloading. This was in the Great Lakes, with seasonal shipping from loading ports to discharge ports. We have a set of loading ports and a set of discharge ports, and each ship takes a full cargo of a single commodity from one loading port to one discharge port. There are many commodities in the problem, but each shipload is a single commodity. You can divide the problem into several problems with separate commodities. The problem is to assign the ships to routes, and there are different types of ships of different sizes. The objective is to minimize the number of ships used. The decision variables were, for each ship, the number of voyages between every pair of loading and discharge ports. They rounded off non-integer solutions, but usually the numbers were in the teens, so I do not know to what extent there was a large error. They used more of the larger ships because they tried to minimize the number of ships. The smaller ships were out of the solution, i.e., they were not used at all. One ship was marginal, used part of the season.

Whiton added some handling capacity constraints to this problem. Rao and Zions, with the same problem, allowed chartering of ships and minimized chartering and operating costs.

I treated a somewhat different problem, a short-term routing problem in which cargoes are assigned to an available fleet in order to minimize the total costs of shipping the cargoes and operating the fleet.

OTHER PROBLEMS

LADERMAN, GLEIBERMAN AND EGAN (1966) - ROUTING PROBLEM, ASSIGN SHIPS TO ROUTES. DIFFERENT SHIPS, SHIPLOADS OF A SINGLE COMMODITY, ONE LOADING AND DISCHARGE PORTS, MIN. NUMBER OF SHIPS. (LP - ROUNDING)

WHITON (1967) - ADDED HANDLING CAPACITY CONSTRAINTS.

RAO AND ZIONTS (1968) - SAME PROBLEM. CAN CHARTER SHIPS. MIN. CHARTERING AND OPERATING COST.

RONEN (1979) - SHORT TERM ROUTING. ASSIGN CARGOES TO AVAILABLE FLEET. SINGLE ORIGIN MIN. TOTAL OPERATING COSTS.

FIGURE 3.10

Lately several interactive scheduling systems have appeared (see Figure 3.11). One is by Stott and Douglas, and Burnie Douglas is here with us today. This deals with a medium-term linear program to assign ships to voyages defined in advance at minimum operating cost. And they have short-term, calculative models, including a simulation model, which are used to compare short-term schedules on a cost basis.

Tom Baker, who is also sitting here, described a problem involving the distribution of petroleum products from a single refinery. There is a linear programming model to select voyages and a network model to determine shipment sizes. It is an interesting extension in that the shipment sizes are not given. They are some of the decision variables.

Let me summarize briefly, using Figure 3.12.

First, the problem of routing and scheduling ships is a complex one.

Second, uncertainty in operations limits the applicability of deterministic models to medium- and long-term planning. In short-term planning, things are always changing; we have to update the schedules very often.

There are several additional variables which I would like to see in models, such as cruising speed. Cruising speed has not been considered in any of the models as a decision variable. For a ship which may be burning tens of thousands of dollars of fuel oil every day, a cruising speed reduction of 20 percent can result in savings of about \$10,000 or \$15,000 a day. Definitely, cruising speed should be part of the decision structure when you are dealing with scheduling ships.

Shipment sizes should also be decision variables, especially when we deal with industrial operators, where shipments are usually on some continuous basis.

Overall, we have the potential for large savings in scheduling ships and routing ships, provided the models are realistic and can be applied to the right problems. This is, in a nutshell, what I have to say today.

INTERACTIVE SCHEDULING SYSTEMS:

STOTT AND DOUGLAS (1981) - MEDIUM TERM LP TO ASSIGN SHIPS TO VOYAGES AT MIN. OPERATING COSTS. SHORT TERM CALCULATIVE MODEL.

BAKER (1981) - DISTRIBUTION OF PETROLEUM PRODUCTS FROM A SINGLE REFINERY. LP TO SELECT VOYAGES AND NETWORK MODEL TO DETERMINE SHIPMENT SIZES.

FIGURE 3.11

SUMMARY

- * PROBLEM IS COMPLEX.
- * UNCERTAINTY LIMITS APPLICABILITY OF DETERMINISTIC MODELS TO MEDIUM AND LONG TERM PLANNING.
- * CRUISING SPEED AND SHIPMENT SIZES AS DECISION VARIABLES.
- * LARGE SAVINGS POTENTIAL PROVIDED MODELS ARE REALISTIC.

FIGURE 3.12

DISCUSSION

MAYS: I find it curious that most of the papers you cited come from the late 1960s and early 1970s. What happened in the latter part of the 1970s?

RONEN: I would guess that people found out that the solutions were not realistic. The more realistic the problem became, as in the paper by Bellmore, Bennington and Lubore where you do not see an example, the harder it was to solve.

Actually, you see now several interactive models, which are returning to the simulative mode, especially for medium- and short-term scheduling.

MAYS: Do you have confidence that we now have the techniques and capability to provide realistic solutions to these problems?

RONEN: No, I do not. The major problem is the uncertainty. The stochastic aspects of the problem are not addressed in the mathematical programming models. That is the major drawback which I see.

WEBSTER: To add to that, I know that in the APL system it is rare for a given ship to keep to its route for more than two or three voyages, because the demand and supply change seasonally. And then there are long-term and cyclical changes always. So it is very hard to find a steady state.

RONEN: It is very hard, definitely. There is no steady state. Things change all the time. And the question is how to address this problem.

WEBSTER: It is a dynamic program, rather than a linear program.

T. BAKER: In the area of using interactive approaches to model building of a stochastic nature, one of the main reasons for going to the interactive type system was to allow the schedulers to do risk aversion, which the scheduler does normally anyway, without having to express it in mathematical modeling form. It gives him a way to bring that into the

problem. And then you do not have to include it.

DOUGLAS: For example, take speed and fuel. That is the reason they do not have to be in the model, because it is interactive and you can inject them.

RONEN: Is this is one of the reasons you went to an interactive model?

DOUGLAS: Yes.

YOUNG: Are there any models that you have come across that have as the objective to make the delivery on time, rather than economic considerations?

RONEN: The earlier tanker models had on-time delivery objectives. The one by Dantzig and Fulkerson and the one by Flood had delivery-on-time objectives as constraints, not as an objective.

YOUNG: And they were trying to maximize the profits?

RONEN: They were trying to minimize the size of the fleet or to minimize the distance in ballast.

In the Bellmore, Bennington, and Lubore model, which is more realistic, you can assign a large utility to delivery on time. And in that case you can transform it into a delivery-on-time objective.

BENTLEY: Related to delivery on time, one of the earlier speakers mentioned inventory carrying cost, time in the pipeline, as a potentially important consideration in looking at the surface effects vessel. You do not list that as an additional dimension to get more reality. Do you think it is not a major consideration?

RONEN: It may be a consideration. But I think speed is much more important.

BENTLEY: On oil cargoes?

BISHOP: With oil, the speed would not even be entered. We determine the speed ahead of time for the year. We found when we did our speed optimization, and we based it on today's market, by the time we got the ships slowed down, the market had changed. And we were always doing the wrong thing; we were sort of chasing ourselves. After a while, we decided to just set the speeds. But the inventory cost on oil at \$34 per barrel at today's value of money is significant.

RONEN: When you talk about speed of a vessel, include the daily cost of the ship; if you own the cargo include the daily cost of the cargo also. That is the inventory-carrying cost there, and that is part of the input to the speed decision.

MACLEAN: We have been discussing this morning the question that you brought up regarding weather routing, as against routing, and as against scheduling. It seems to me that we would save some confusion if we recognized that, when you said you were going to concentrate on routing, that really had to do with trade route definition and vessel assignment.

This becomes one order of the hierarchy. Below this, then, comes the vessel scheduling. And below that, something that nobody has mentioned this morning -- passage planning - and a subset of that becomes the direct vessel routing. If we can keep that kind of hierarchy in mind, I think it will help our discussions.

I'd like to say something more about what Bill Webster has indicated; quite clearly, a primary function of any organization today is to maximize its return on invested assets. That is, first and foremost, essential to survival in an inflationary economy, because if you cannot maximize greater than the inflation rate, you are going to fail. Within that constraint, one then has a whole hierarchy of operational factors that are involved. Once you are capital-invested in a particular trade, then you can degenerate to a maximization of profits. You can go further on down to where you want to maximize for a given passage. You want to maximize net revenues, income over outgo. And you can keep that hierarchy in a systematic structure. I think it would be helpful,

in our discussions today and tomorrow, if we try to keep this hierarchy in our thoughts as we talk about it.

WILMES: Mr. Young asked you a question about whether or not you have come across, in your literature research, any schedulers to meet delivery dates. The government has one. It is called SEACOP. SEACOP is ten years old. It is based on the state of the art 20 years ago.

We are concerned, in our new SEASTRAT modeling, about doing what was alluded to by Mr. Young. In looking at a couple of the other papers, one of the things that we are concerned about, particularly in the definition of "strategic mobility," has to do with meeting the vital dates of the combat commanders. It goes back to the question of meeting delivery dates. We do have a model, but it is not good. So this is one area where we would like to draw on your talent and the experience of other people in this room.

RONEN: Is the SEACOP model published?

WILMES: It is not published.

RONEN: I confined myself to published sources.

SOLAND: I would like to echo Colonel Wilmes's request. I think we could all benefit from a digging for further references. Many of you may know of additional papers that are appropriate to list along with the ones that David has referenced. You should all have a copy of his paper. I know of a couple of others that do not appear there. They may appear elsewhere in his work.

I would really appreciate it, and I think other people here would appreciate it too, if you would search your minds a little bit and write down any references that you have on the kinds of problems that we have dealt with and are dealing with, whether or not they are in the open literature. I am sure that there are some private company reports that may be available to the general public and some Department of Defense reports which are nonclassified and are available. Let us get these

listed down for everybody's benefit, so that we can know what problems have been attacked, what techniques have been used, et cetera. We will try to accumulate these so that we all will benefit from a comprehensive set of references.

The Utility of Heuristics in
Relatively Complex Strategic Mobility Models

Carroll J. Keyfauver

General Research Corporation
McLean, VA

I will start with an outline (see Figure 4.1) of the kinds of things I will talk about; first, some of the elements of strategic mobility, and specifically the deployment scheduling problem with which we are mostly concerned. Then I will talk about some of the problem areas in modeling strategic mobility and an approach to the problem, which takes into consideration some of the dynamics of the problem, some of the heuristics, and finally some of the problems we encounter using heuristics.

Now, here is a definition. "Strategic mobility" is the capability to move military forces and supplies to where they are required, during a period of actual or potential conflict. There are a number of different ways in which the strategic mobility problem arises. One is the contingency planning problem that was described earlier, summarized in Figure 4.2 as the deployment scheduling problem. The other problem is in doing analysis concerned with the resources that are needed for strategic mobility planning, particularly, in future years, and analyzing strategic mobility programs. My experience and background is mostly in that area, but a lot of it carries over into strategic mobility scheduling for contingencies.

I have placed the elements of strategic mobility into three different areas (see Figure 4.3).

One is the cargo that we have to move--unit equipment, personnel, and supplies.

FIGURE 4.1

- ELEMENTS OF STRATEGIC MOBILITY
- THE DEPLOYMENT SCHEDULING PROBLEM
- PROBLEM AREAS IN MODELING
- A SEQUENTIAL ADAPTIVE APPROACH
- A HEURISTIC SELECTION RULE FOR CARGO SCHEDULING
- PROBLEM AREAS IN USING HEURISTICS

THE DEPLOYMENT SCHEDULING PROBLEM

- FIND A SCHEDULE OF SHIPS AND AIRCRAFT TO MOVE UNITS AND SUPPLIES WHICH SATISFIES OR MOST NEARLY SATISFIES THE OBJECTIVES OF THE DEPLOYMENT.

FIGURE 4.2

GRG [R]

ELEMENTS OF STRATEGIC MOBILITY

- CARGO
UNIT EQUIPMENT
PERSONNEL
SUPPLIES
- TRANSPORTATION ASSETS
AIRCRAFT
SHIPS
PREPOSITIONED EQUIPMENT AND SUPPLIES
FORWARD BASES
PORT FACILITIES
CARGO HANDLING EQUIPMENT
- SCENARIO
RESOURCE MOBILIZATION
ROUTES FOR AIRCRAFT AND SHIPS
ATTRITION OF AIRCRAFT AND SHIPS
CONVOYING OF SHIPS

 GRC

FIGURE 4.3

Transportation assets, which are ships, aircraft, prepositioned equipment and supplies, forward bases, port facilities, and cargo handling equipment form the second area.

Then there are the elements that describe the scenario in which the deployment will occur--how the resources are mobilized, what routes will be used for aircraft-to-ship operations, whether there will be attrition, what kind of attrition there might be, and whether or not ships convoy. These are examples of some of the things that are dependent upon the scenario.

For the cargo (see Figure 4.4), we can assume that we have a current known location, a known readiness, a pre-assigned destination, and a required delivery date, a date that is specified by the theater commander when he would like to have the cargo reach the destination. We also have a known composition of equipment and supplies. We know how many pieces of equipment and what kind of equipment there is. We also may know whether it is required to be air lifted. If it is large, it cannot be air lifted; so, therefore it has to be sea lifted. We know whether, in some cases, there is a designated port through which the cargo would have to be moved.

For aircraft and ships (see Figure 4.5), we generally have data on what the current location may be, either as a statistical distribution, based on some type of snapshot, or current data supplied from the ships and aircraft. Also, we have an expected payload. In general, we have some idea of how much of each type of cargo the ship or aircraft might be expected to carry. We also have known operating characteristics of the ship--speed, range, draft, and information on cargo-handling, whether or not it is a container ship, a break-bulk ship, a RO-RO, etc.

For the scenario (see Figure 4.6), we know the locations of the bases and ports we will be using. We generally have a good idea of the transportation route network that we will be moving over. We also have some idea of the expected requirements for supply. We also have a known supply policy, that is, how much will be required each day, how the

CARGO

1. KNOWN CURRENT LOCATION
2. KNOWN READINESS
3. ASSIGNED DESTINATION
4. REQUIRED DELIVERY DATE (RDD)
5. KNOWN COMPOSITION
6. DESIGNATED MODE?
7. DESIGNATED PORT?

AIRCRAFT AND SHIPS

- 1. CURRENT LOCATION**
- 2. EXPECTED PAYLOAD FOR EACH CARGO TYPE**
- 3. OPERATING CHARACTERISTICS**
 - a. SPEED**
 - b. RANGE**
 - c. DRAFT**
 - d. CARGO HANDLING**

FIGURE 4.5

GR(4) [4]

SCENARIO

- 1. LOCATIONS OF BASES AND PORTS**
- 2. TRANSPORTATION ROUTE NETWORK**
- 3. EXPECTED CONSUMPTION RATES FOR SUPPLY**
- 4. SUPPLY POLICY**
- 5. MOBILIZATION DATES**
- 6. ATTRITION RATES**
- 7. CONVOY POLICY**

FIGURE 4.6

AD-A123 464

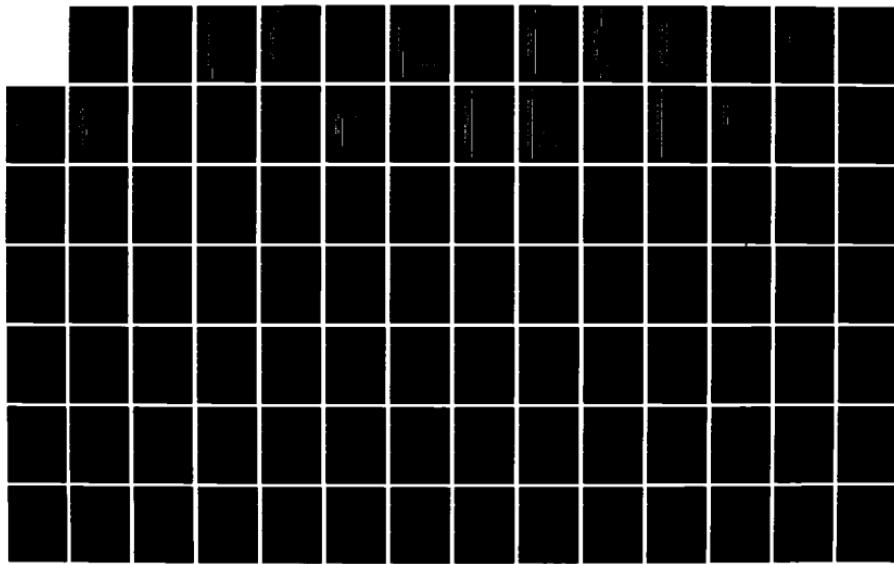
PROCEEDINGS OF A SYMPOSIUM ON CARGO SHIP ROUTING AND
SCHEDULING HELD WASH. (U) GEORGE WASHINGTON UNIV
WASHINGTON DC INST FOR MANAGEMENT SCI. R M SOLAND

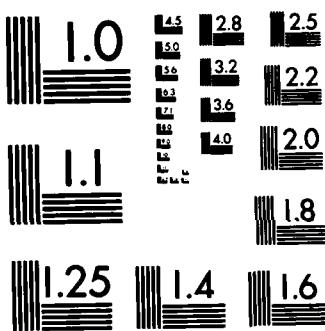
2/4

UNCLASSIFIED

15 DEC 82 SERIAL-TM-69250 N00014-82-G-0016 F/G 15/5

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

supplies will be stockpiled, where the supplies are, and so forth. We also have mobilization dates for both the cargo and for the transportation resources. We also, in general, may have some estimates of attrition rates and what the convoy operations policy would be for a particular type of contingency.

There are four problem areas in modeling strategic mobility that I am going to discuss particularly in reference to the deployment scheduling problem (see Figure 4.7).

First is objectives and objective functions; second is the size of the problem; third is the dynamic nature of the problem; and the fourth are the uncertainty and probabilistic elements.

See Figure 4.8. There are at least two major approaches to trying to define objective functions for strategic mobility planning.

One is to minimize costs for transportation resources, and generally this is the one that arises in trying to analyze new strategic mobility programs--to decide whether or not we should buy new aircraft or new ships, or prepositioned equipment, or any of a number of programs. The problem with using minimum cost in a contingency planning, operation planning, or execution situation is that you have no time to buy new resources--new ships or new aircraft. They are very long lead-time items and you are generally stuck with what is available.

Another characteristic of the minimum cost problem is that, in general, the solution is very sensitive to the statement of the required delivery dates (RDDs) of the cargo, particularly early in the deployment. A very small variation in RDDs can cost you billions of dollars in additional resources, just to get a small increment in capability quickly.

The other objective that has been used is to maximize capability. There are a number of problems with maximizing capability, however. One is to describe exactly what you are maximizing. Maximizing the capability of the transportation resources available to you may not maximize the capability of the forces that you can deploy. It is very

PROBLEM AREAS IN MODELING

- 1. OBJECTIVES AND OBJECTIVE FUNCTIONS**
- 2. SIZE OF THE PROBLEM**
- 3. DYNAMIC NATURE OF THE PROBLEM**
- 4. UNCERTAINTY AND PROBABILISTIC ELEMENTS**

FIGURE 4.7

OBJECTIVES AND OBJECTIVE FUNCTIONS

- MINIMUM COSTS
FIXED TRANSPORTATION RESOURCES
VERY SENSITIVE TO RDDs
- MAXIMUM CAPABILITY
MAY NOT HAVE FEASIBLE SOLUTION
TRADEOFF OF HIGHLY INTERRELATED
BUT DISSIMILAR ITEMS
VALUE OF TIME

GRG C

FIGURE 4.8

difficult to precisely define the relationships among the kinds of cargo we are moving, between unit equipment and supply, and POL, and the various categories of supply, such as ammunition and food.

All of these things are needed in order to render functional a combat unit that is being deployed. Yet, once you have attained a certain supply level, they really add no additional capability to the force that you are deploying.

And there is no guarantee that there is any feasible solution to the problem, i.e., there is no reason to expect that there is a feasible solution that meets all of the RDDs.

Also, there is the problem that, given two units that deploy, one early in the deployment and the other one much later, a delay with respect to the early unit may be much more significant and affect the outcome of the deployment much more greatly than a delay with respect to the deployment of the later unit. Later in the deployment when you have more forces, however, you have more alternatives and a delay is generally less critical.

Let us talk about the size of the problem. I think Figure 4.9 gives an idea of the size of a deployment. We are talking about possibly 10,000 units. We are talking about forces that may range from the size of an armored division down to two-man units that perform some special mission, say the two-man well-drilling team that shows up frequently in a number of the deployment problems.

When you look at it in terms of pieces of equipment, there are literally millions to be deployed--of all sizes, ranging from tanks to rifles to computer test equipment. And at the same time, we are deploying hundreds of thousands of people to man all of these units.

Also, we may have on the order of 1,000 aircraft that could be used in a deployment situation, and perhaps 1,000 ships, for 180 days.

Also, you have the possibility of alternative ports to which cargoes could be scheduled, and you have the possibility that within a port, alternative port facilities could be available for each unit that is being deployed.

SIZE OF THE PROBLEMS

- 10000 UNITS
ARMORED DIVISION AND SMALLER
MILLIONS OF PIECES OF EQUIPMENT
HUNDREDS OF THOUSANDS OF
PERSONNEL
- 1000 AIRCRAFT
- 1000 OR MORE SHIPS
- 180 DAYS
- ALTERNATIVE PORTS
- ALTERNATIVE PORT FACILITIES

FIGURE 4.9

The dynamic nature of the problem is one of the areas that we had the most difficulty in dealing with (see Figure 4.10). Either side may escalate the conflict. The problem may change at any time. There is no way to lay out a deployment problem and guarantee that a deployment schedule that you have today will be the deployment schedule you need tomorrow. The problem may be completely different. You may be deploying in a different direction, with different forces and with different resources. The new deployment problem may require re-allocation of the resources and it may also require rerouting of cargo.

I would like to mention some of the kinds of probabilistic events that can occur during a deployment (see Figure 4.11). We can have attrition of aircraft; we can have attrition of ships; we can have losses of ports or port facilities. There is an almost endless number of things that can go wrong.

In general, if you have a very detailed schedule, you can expect that it will not hold up for very long. There is also a limit to what you are capable of planning and controlling, simply because of the huge number of pieces of equipment that you are deploying.

There are numerous opportunities for over-estimating or under-estimating the time it will take to perform a particular activity, and you simply cannot take an individual box and track it all the way through from its depot through each stage--onto the truck, off at the port, onto the ship, and through each stage to the destination. It simply serves no purpose to plan at that level of detail; the return on investment in planning would be virtually nil.

What I would like to describe next is an approach for gaining some control over this problem (see Figure 4.12). Essentially, what it involves are the following steps.

First you develop a method for solving a deployment scheduling problem.

Then you perform the operations that you derive from that schedule for a period of time.

DYNAMIC NATURE OF THE PROBLEM

- 1. EITHER SIDE MAY ESCALATE THE
CONFLICT**
- 2. REALLOCATION OF RESOURCES**
- 3. REROUTING OF EN ROUTE
CARGO**



FIGURE 4.10

UNCERTAINTY AND PROBABILISTIC EVENTS

- ATTRITION OF AIRCRAFT AND SHIPS
- INTERDICTION OF PORTS
- LIMITS TO PLANNING AND CONTROL

A SEQUENTIAL ADAPTIVE PROCESS FOR DEPLOYMENT

- STEP 1 DEVELOP A DEPLOYMENT PLAN BY
 SOLVING THE DEPLOYMENT
 SCHEDULING PROBLEM**
- STEP 2 EXECUTE THE PLAN FOR A PERIOD OF
 TIME**
- STEP 3 REASSESS THE STATE OF THE
 DEPLOYMENT AT THE END OF THE
 PERIOD**
- STEP 4 DEVELOP A NEW DEPLOYMENT
 SCHEDULING PROBLEM**
- STEP 5 GO TO STEP 1**

GRGC

FIGURE 4.12

Then you reassess the state of the deployment--the location of all of your resources, the location of all of the cargo to be deployed--and you take into account any events that have occurred which would change the deployment or change its direction, e.g., an expansion of the war which would cause a redirection of your planning.

Then you develop a new deployment scheduling problem, taking the state of the system where you stopped as the initial conditions from which you solve the new problem.

Then you repeat this process over a period of time, iteratively, for whatever period of time over which you want to do your deployment planning.

What are some of the benefits of this kind of approach (see Figure 4.13). One is that it adapts the solution to the deployment scheduling problem to the dynamics of the situation. And it allows you to adjust the deployment schedule to take into account unforeseen events.

Also it is equally useful for the planning or the actual execution of deployment. There is no restriction as to how this method can be applied. In an actual deployment situation, you could use real data which is fed back from the ports and from the units that are being deployed. For planning, you can use a simulation model to do the detailed modeling that would take into account a lot of the things that you aggregate in the deployment scheduling process.

DOUGLAS: Back on Figure 4.12, do I understand that Step 2--execute the plan for a period of time--is not as long as the deployment plan overall? You just execute part of it and see how it goes?

KEYFAUVER: Right. We talk about executing it for a small increment of time, maybe a day, maybe an hour. I have not specified the period of time here.

We have built some models along these lines, and at the moment we are using one day as the increment of time. For sealift, that is not a long period of time; for airlift operations, it is a very long period of time.

BENEFITS OF THE SEQUENTIAL ADAPTIVE APPROACH

- 1. ADAPTS SOLUTIONS TO DYNAMIC SITUATIONS**
- 2. ALLOWS ADJUSTMENT OF THE DEPLOYMENT SCHEDULE**
- 3. EQUALLY USEFUL FOR PLANNING OR ACTUAL EXECUTION OF A DEPLOYMENT**

FOR PLANNING, A SIMULATION CAN BE USED TO DYNAMICALLY TRACK EVENTS

FOR OPERATIONS, REAL WORLD EVENTS CAN BE USED

PROBLEMS

- 1. MANY MORE SCHEDULING PROBLEMS TO SOLVE**

FIGURE 4.13



The problem with this approach is that it does not solve our scheduling problem. We do not have a solution, but instead we now have a whole series of problems, instead of just one problem.

Now we will talk about approaches to dealing with the deployment scheduling problem. Figure 4.14 is an example of what may happen to two identical units. They do not even have to be identical, but for my purposes, suppose we have two units which have the same RDD at their destinations; they are going in geographically different directions. One is going to Europe, and the other to the Persian Gulf. One has an expected travel time of 15 days, while the other has an expected travel time of 35 days for the purposes of illustration ("C" stands for deployment date).

What do we expect as the latest departure date? When will we expect to have to start moving each unit in order for it to arrive at its destination on schedule? Obviously, what this is saying is that we have to be concerned with deploying to the more distant destination much earlier. The Persian Gulf unit has a high priority requirement for resources, it needs to be scheduled much earlier in the sequence than a unit which is going to a much closer destination.

I will describe one heuristic selection rule which we have used for selecting a sequence of cargoes to schedule in our model (see Figure 4.15).

Basically, the rule is to select in order of priority, determined by the cargo's RDD minus the expected travel time from origin to destination through the port which gives the shortest total expected travel time, taking into consideration land movement on both ends and the time at sea. There may be several ports through which the unit could be moved. You may have a number of possibilities for the deployment route, and different times at which the unit would be required to depart its initial location in order to meet the schedule.

EXAMPLE

**IDENTICAL CARGOES
FROM EAST COAST TO:**

PERSIAN GULF

EUROPE

C+45

C+45

35 DAYS

15 DAYS

**EXPECTED TRAVEL
TIME**

C+10

C+30

**LATEST DEPARTURE
DATE**

GRC

FIGURE 4.14

A HEURISTIC SELECTION RULE FOR CARGO SCHEDULING

- SELECT CARGOES TO SCHEDULE ONE AT A TIME IN ORDER OF PRIORITY DETERMINED BY THE CARGO'S RDD MINUS THE EXPECTED TRAVEL TIME FROM ORIGIN TO DESTINATION THROUGH THE PORT OF EMBARKATION WITH SHORTEST EXPECTED TRAVEL TIME.

FIGURE 4.15

GRC

We have chosen the shortest expected time because, intuitively, it tends to minimize the amount of transportation resource required to move that unit.

DOUGLAS: Could you spend a little time on what makes something heuristic versus whatever else there is?

KEYFAUVER: I guess one definition of "heuristic" would be any method which leads to a solution, but one that cannot be guaranteed to be an optimal solution.

DOUGLAS: Is it trial and error?

KEYFAUVER: Not necessarily. It may be a systematic set of rules that leads you to a solution that you feel is in the direction in which you want to go. It is better than trial and error.

DOUGLAS: Try and try again. Maybe that is a better definition. I am not trying to pick words. I just have not seen any difference yet.

MAYS: Have we determined that there is a need for the heuristic selection rule and that we cannot use standard techniques?

KEYFAUVER: What is a standard technique?

MAYS: Can we define what the problem is, so that we can see if we want to give up on some of the mathematical tools that are available?

KEYFAUVER: We are talking about 1000 unique ships, 10,000 units and 180 days; and when you start multiplying some of the possibilities, we are talking about a problem of very large size with a large amount of computation, especially if you get into some of the things like queueing where you actually have a multiplication of time periods. You may have to take into account the queueing at destination ports, and so if you use something like LP you have a situation where you have got to schedule the unit to arrive when there is no queue, or there is a queue of one day, or there is a queue of two days, and so forth.

TEMMLER: Are you saying that there are restraints on the amount of time that you want to give to solving the problem, is that what you're saying, or cost constraints?

ENGLISH: Time is the military consideration, if you must think of one constraint. I am sure we can persuade you that time is a vital constraint as far as the military problem is concerned, and it is that which brings a number of us here to this meeting.

TEMMLER: That is not my question. Jim Mays asked why go heuristic rather than to an absolute method.

ENGLISH: Since I provoked Carroll into giving this presentation I would like to comment. Carroll represents an organization that has a mighty stature in this general endeavor. They have built some models and, although Carroll was explicitly asked not to cite specific models that he has built in DoD, they are real triumphs as far as that sort of thing goes.

I believe you can take it as part of the reason for which we are here, that when he and his group undertook this effort they concluded that they could not handle the scope of their problem analytically. The complexity of their problem could not be handled any other way, and that, in part, is one of the things he is trying to convey. The hope is that this will bring us to one of the issues with regard to the next family of models, particularly the one with the Military Sealift Command: how they ought to go, whether they should be entirely heuristic or still seek to work with some of the more traditional mathematical solutions.

TEMMLER: The question still revolves around the words, "They concluded that they could not use other alternatives." Did they conclude that because they felt it would take 15 of them 10 years and it was too costly?

KEYFAUVER: Let me try to answer that. We could spend several hours talking about that problem. There are a number of optimization techniques which could be used. The problem is that it is not clear that there are any that are capable of handling a problem of the scope and size that we are talking about.

We are talking about a problem of such size that if you were to define it in the form of, possibly, a binary or linear program, it would probably run into the hundreds of thousands, or millions, of rows as constraints. And I would not want to guess how many columns it would have.

There are a number of other techniques you could approach it with, including various network analysis techniques. Many of them are limited to solving smaller problems and using heuristics as a part of the approach. But there are techniques that I do not know about. There is one of the things I am here to learn about.

Let us go on to Figure 4.16. All we have accomplished is simply to reduce the size of the problem. The way I wrote the heuristic selection rule, it implies that these requirements are processed one at a time. That is not necessarily the way in which you would implement it. You could use the same rules to process all the requirements that come up with equal priority, or you could select within a range of units those that fall within a certain time span, or certain priority range, for your consideration as a set of problems to solve simultaneously, using an optimization of some type.

Presently, the technique I am describing is solving problems one at a time, and reaching an optimal solution for that single problem, by implicit enumeration of the alternatives, by looking at all of the alternatives that you can have for that single scheduling problem and trying and choosing the best of those alternatives.

As I said, Figure 4.16 points out that we do reduce the deployment scheduling problem to a sequence of much smaller problems, and the other benefit that we gain is that only some of the problems have to be solved

BENEFITS OF THE SELECTION RULE

- 1. REDUCES EACH DEPLOYMENT SCHEDULING PROBLEM TO A SEQUENCE OF MUCH SMALLER PROBLEMS**
- 1. ONLY SOME OF THE PROBLEMS NEED TO BE SOLVED AT ANY ONE TIME**
- 1. ONCE ALL SHIPS HAVE BEEN SCHEDULED, WE CAN STOP.**

FIGURE 4.16

GRCA

at any one time. Once we have scheduled the entire set of resources available to us at any one time, if we have scheduled all of the ships that we have available or will have available for some foreseeable future time, including all of the ships that are just leaving a port after having made a delivery, then we really do not obtain any benefit from continuing this process indefinitely. We are able to describe the projected state of everything in the system, and we can go and execute this plan for some period of time, and then do a new evaluation of the state of the system by developing a new deployment scheduling problem.

Some of the problems that we have with this kind of selection rule are shown in Figure 4.17. There is a cargo availability problem. Also, there is the problem of not being able to determine the quality of the solution obtained by using heuristics. We have no way of measuring how close we are to an optimal solution using this kind of approach. That is a serious problem.

There is, also, what I call the "ship circling" problem, the "closest port" problem, and the "supply" problem. Look at Figure 4.18.

The cargo availability problem is simply that if you take the cargoes in order of the priority selection rule that we have used, and just assign the first available ship from among those that are in port, some of those ships may wait around a long time for their cargo to arrive at that port, which is not an efficient utilization of those shipping resources.

Conversely, in trying to schedule a unit, you have a critical path situation where the last cargo to arrive at the destination determines the closure date of that unit, and, thus, when that unit becomes effective. Therefore, the ships that arrive early really contribute nothing until that unit becomes a whole entity at the other end.

A solution that you could use is essentially another heuristic approach: do not have ships wait unless the cargo is about to be at the port. This is a rather loose statement. Rather than giving precise numbers for how long you should wait, I want to describe the kinds of problems that you have to deal with in using this approach.

PROBLEMS WITH THIS KIND OF SELECTION RULE

- 1. CARGO AVAILABILITY PROBLEM**
- 2. NO WAY TO DETERMINE THE QUALITY
OF THE SOLUTION**
- 3. SHIP CIRCLING PROBLEM**
- 4. CLOSEST PORT PROBLEM**
- 5. SUPPLY PROBLEM**

FIGURE 4.17

GRG

CARGO AVAILABILITY PROBLEM

- **PROBLEM**
**SHIPS WILL WAIT FOR CARGO TO
ARRIVE AT PORT WHEN LOWER
PRIORITY CARGO IS AVAILABLE**

- **POSSIBLE SOLUTION**
**DO NOT ALLOW SHIPS TO WAIT FOR
A CARGO UNTIL THE CARGO IS AT
THE PORT**



FIGURE 4.18

Using this kind of heuristic creates a whole new set of problems that you would not get if you could use one of the optimization techniques, particularly LP, where you can explicitly enumerate all the solutions, and the algorithm itself will wash out the solutions that make very little or no sense.

The ship circling problem (see Figure 4.19) follows from the question: to which port do you assign the ship after it has delivered the cargo? If you use successive solutions and a sequential process such as I have described, you have the possibility that the next time you generate a new schedule the ship is assigned to go in the opposite direction from the one it was sailing in or, in effect, it could end up just sailing in circles. So you have the problem of how to avoid that, how to cause the ship to keep going in whatever direction it should continue to go.

One approach you can take is to recognize that this ship is headed in a certain direction, and that it is moving away from port, and to simply not allow it to turn around and go back unless there is some overriding consideration.

There is also the closest port problem (see Figure 4.20), which is a geographical phenomenon. For example, suppose a ship is returning from Europe and going to a port on the Gulf Coast. It passes very near all of the ports on the East Coast, and allowing the model to go on its own, the model would show a tendency to grab that ship and put it into the nearest East Coast port because it is there first, or it could potentially be there first, even though the cargo that you are thereby selecting may have a lower priority.

There are some solutions for this type of a problem, e.g., recognizing ships further from the port and giving the more distant ports a headstart in selecting the resources they need, because they have the more serious scheduling problems.

DOUGLAS: Would that depend on the priority of the cargo though, rather than just on the port?

SHIP CIRCLING PROBLEM

- PROBLEM
 - SHIPS MUST BE ASSIGNED A PORT TO WHICH TO RETURN AFTER DELIVERING A CARGO. SUCCESSIVE PROBLEM SOLUTIONS COULD CAUSE THE SHIP TO SAIL IN CIRCLES.
- POSSIBLE SOLUTION
 - THE SHIP WILL BE ONE DAY CLOSER TO THE ASSIGNED PORT ONE DAY AFTER THE PORT IS ASSIGNED. PREVENT THE SHIP FROM BEING REDIRECTED TO A PORT IT IS MOVING AWAY FROM.

CLOSEST PORT PROBLEM

- PROBLEM
 - WHEN A SHIP MUST SAIL PAST ONE PORT TO ARRIVE AT ANOTHER TO LOAD CARGO, THE SHIP MAY BE ASSIGNED TO THE CLOSER PORT EVEN IF THE OTHER PORT HAS PRIORITY CARGO.
- POSSIBLE SOLUTION
 - ALLOW MORE DISTANT PORTS TO SCHEDULE SHIPS EARLIER

KEYFAUVER: Yes, it depends on the priority of the cargo.

BALLOU: Going back to priority in surface sea transportation, can you give an example, say, where priorities would not be shown by, say, the RDD or the arrival date? I have even run into a prioritized list. I have never heard of anyone assigning it for sea.

KEYFAUVER: The priority I am talking about is based in part on the RDD and, taking into consideration the expected or the estimated transportation time through the system, trying to determine the sequence in which to mobilize cargo and haul it to the port for deployment, within CONUS, so that we can get it to the destinations on time. There is a definite sequence that takes into consideration the fact that some cargoes have got to go a lot further than others.

Another problem of concern, regardless of the approach you use, is estimating your supply requirements (see Figure 4.21). Suppose you use just the RDD as an estimate of your supply requirement; initially, that is fine. But if you manage to get ahead of your deployment schedule, or get behind, your requirements may deviate substantially from that original estimate. And, there is another consideration, too, which is the actual combat that may occur in a theater. It may be at a substantially different level than expected. Maybe you were planning on the war breaking out and it did not. So you have a reserve supply that you have not used up, that you were not anticipating.

This sort of thing can be taken into consideration, because what we have is a process that looks ahead, forecasts the arrivals of units, allows us to estimate what the supplies would be at some future date, and, therefore, allows us to anticipate the need for those supplies and to ship them in anticipation of those needs.

SUPPLY PROBLEM

- PROBLEM
 - SUPPLY REQUIREMENTS MAY BE SUBSTANTIALLY OVER ESTIMATED OR UNDERESTIMATED BY USING UNIT RDDs
- POSSIBLE SOLUTION
 - SUPPLY REQUIREMENTS CAN BE FORECAST FROM THE SOLUTIONS TO EACH DEPLOYMENT SCHEDULING PROBLEM

FIGURE 4.21

GRG

DISCUSSION

MACLEAN: It seems to me that this is not unlike a commercial problem in the sense that before you can arrive at a decision on the relative importance of movements you have to be able to quantify in a consistent fashion the value of getting something from point A to point B. If gasoline to supply a truck in support of a battalion movement in a particular advanced base is worth a lot to you, this is the driving function that makes you decide to send a ship from point C to point D to get over there and get that thing delivered in time. You have got to have a driving function that is consistently quantified through your whole structure, or regardless of how heuristic or optimal your approach is going to be, you are going to end up with a mish-mash of inconsistencies that constantly creates frustration. You have got to get to the quantification of the issue. Otherwise, you cannot make consistent decisions.

KEYFAUVER: I completely agree with you.

MACLEAN: Where does this come into your approach here?

KEYFAUVER: It comes in in the sense that there is an implied objective in this process, which is to try to meet the RDD, to minimize the lateness of the arrival of the unit, and to try to get the sequence of units there as early as possible. As you revise and develop a state of the deployment each day, you know when a unit is needed, you know where it is presently, and you can develop some idea of how long it will take to get it from where it is to where it is needed. And that determines a sequence, which, in effect, is used to determine the priority that I am talking about.

MACLEAN: You have an inherent conflict there that you cannot solve this way, because one RDD has no relationship whatsoever to another RDD.

KEYFAUVER: One hopes that they do.

MACLEAN: Then how are you going to quantify this?

HALL: Theoretically, that is done during the planning process by the supported commander's staff. I cannot detail how they do that, but I think the key here is that the required delivery date that the model deals with then theoretically reflects the priority that the supported commander has set. It is decided, unit by unit, by what he needs in a given location by a certain time to execute his combat plan. That is really the genesis of the required delivery date.

MACLEAN: Has he also quantified the consequences of default?

HALL: No.

MACLEAN: Then suppose you do not meet the required delivery date; what happens? What is the consequence? Does everything fall on its face, or is it a minor inconvenience? And if so, if it is only a minor inconvenience, why have you gone to all of this structure to say that is a required delivery date when it does not really mean anything?

HALL: I do not think I would go so far as to say it does not mean anything. There are elements in it that are not quantified. I will admit that much.

MACLEAN: That gets to the root of his problem, it seems to me, in terms of how he decides whether or not he is going to have a ship circling in the ocean, doing nothing, because it has got an unweighted set of RDDs trying to drive it from time period to time period to time period.

KEYFAUVER: They are not unweighted; that is the whole point.

MACLEAN: Well, you have not said anything about weighting yet.

KEYFAUVER: In effect, an RDD is a weighting of a set of units to be deployed.

VOICE: This comes back to the question of the purpose of the analytical treatment. Are two commodities, two pieces of cargo with the same RDD, given equal analytical treatment, all other things like transportability aside?

They are treated equally because there is no other indicator, which I think is what Joe Ballou was trying to get at. There is no other measure of priority aside from RDD. Given all of this other background, e.g., things like same distance to go, then two cargoes with the same RDD are treated, at least in some heuristic approaches, as coequals.

TEMMLER: You are talking about the priority of cargo. That does not enter into it. If two RDDs are equal, the cargo priority does not have anything to do with it.

KEYFAUVER: RDD is a major driving force for the deployment of any kind of cargo. And given that we have assigned RDDs of, say, 25 and 30, we expect the one that has an RDD of 25 to be given a higher priority, that is, earlier consideration than the one that has an RDD of 30, particularly if they are going to the same place.

FOARD: Let me address the issue of the RDD, the CINC, and what happens if a delivery does not meet the RDD. Ideally, the CINC sets the RDD on it when he originally starts, and this is in the real world of MSC, as opposed to where Carroll Keyfauver works and I work, and where one looks at mobility issues such as do we have enough mobility. The RDD is assigned two different ways. The CINC is assigned the forces by the Joint Chiefs of Staff, he is told this is his area of responsibility, and he will allocate those forces, plan on using those forces, and develop a plan to use. He goes through this process and says: "To minimize my risk, in fighting this war against the threat, I would like to have these guys here by this time." Ideally, he would like to have them in place on day zero when he first starts, wherever he is fighting. If he does not get that, they will have to move. He will try to logically sequence them into where he wants, hopefully getting them, and then he runs through a process of constraining that delivery profile, because he never has enough assets to get them exactly when he wants. He sends his plan back here.

And as MSC indicated this morning, we go through an interim process back here with the Joint Deployment Agency, with the JCS, with the TRA and MSC, who run it through their systems, and they refine it down some more. Then we get a profile of what it is going to look like. And then it is subjected to judgment. Is it feasible, is it going to work, or do we have to can the whole thing and start over again? In the real world that is what we are talking about. And that is when we at the JCS level make the determination. Is the plan feasible for logistic mobility, are we getting it there in time? And CINC makes a determination at that point. When he sees when his forces are going to get there, he knows what kind of shape he is in, to evaluate against the threat.

In the studies, we are trying to answer the question: what do we need to build? For example, do we need to build a C-17? We have worked in studies for the out years. And here are RDDs defined by the services. The Army, Navy, and Air Force are looking at minimum risk situations. They say, "You get it here, and we will minimize the risk if you can meet the RDDs." We run that through, and we say, "Here is where we think we can get it to, using the assets of 1986-1987. We have a shortfall here, and if you buy me more airplanes, more fast ships, preposition stuff in the Indian Ocean for one theater or put more into Europe, preposition there for the other theater, then I can shorten that gap." And that is an out-year look at it.

VOICE: It is really an iterative process. The judgment has been made when the RDDs are established that the whole arena has been addressed in establishing the RDDs. That is an input to this problem. The military consequences of not meeting those RDDs are analyzed outside. That will tell us what the delivery schedule is that we can meet.

JESMER: There is one more step that we can take. We can feed the output from our mobility models into a theater war game, which is loaded with the enemy threat and the friendly forces, and you can take a look at the impacts that might result from your prioritization of your cargo movements, as well as what happens when things are

delayed. What is the impact? What might be the impact on the war in that field, as established by the war game?

We do go through that process with the plans, so we have some idea of whether or not we may want to change our priorities, given our lift constraints and the cargo movement requirements as prioritized by the CINC. And many times they will come back and change and reprioritize.

KEYFAUVER: That sort of thing can be done as part of the assessment of the state of the system, but it is generally not a mobility consideration. We have not reached the stage where we can take that into consideration dynamically during a deployment simulation or deployment modeling effort.

JESMER: In actual execution, the Joint Chiefs of Staff have a system for actual cargo movement, and if the CINC determines there is going to be an impact based on some information that he receives about a cargo movement, he can always change the priority. And, of course, in actual execution the cargo would be moved out if it is needed sooner.

YOUNG: How well does the heuristic work? Most of your problems seem to be a consequence of the heuristic itself. How well does it work when you try simple problems where you need a good solution?

KEYFAUVER: So far as we can tell, it works pretty well. One of the problems with this kind of heuristic method is that there is generally not a problem to which you have an analytical solution that will allow you to make that kind of comparison.

YOUNG: Have you ever tried a small problem just to see how well it works in terms of knowing what the best solution is?

KEYFAUVER: The small problems could probably be done by hand. There are not many models that would take into account the kind of detail that this model can.

YOUNG: Have you made any comparisons with the hand schedules?

KEYFAUVER: We have done a little of that. It seems to check out very well. That is just one approach. I am not trying to say that this is the only approach to the problem. There are lots of heuristics. We have only tried one, essentially. There are a few possible variations upon it. I think it is an area that needs a lot more exploration; and, as I said before, at the point at which we are using heuristics, if we had some other technique, for example, an optimization technique, that we could use, fine. We could plug it right in here instead, because it would get us one step further.

YOUNG: Basically, your experience relative to this type of problem is restricted to this one heuristic in trying to fix it up and make it work and not trying to generalize at this point--or are you?

KEYFAUVER: There are other heuristics that we have tried in other models. We seem to be fairly consistent among them in the kinds of results that we achieve.

MACLEAN: One of the things I do not really quite understand about this yet is that I do not see how Murphy's Law gets into this process, because the probability of failure is there in all cases; and if the RDD is an absolute, then you have got to have redundant supply built in. And we have not said anything about that. It seems to me that somewhere along the line we have got to bring in the probability of a proposed action being successfully carried out in order to make sure that the heuristic solution to the problem has a satisfactory probability of coming to fruition.

KEYFAUVER: That is something I was trying to get at. I talked more about the scheduling part of this thing, and I mentioned simulation. The simulation can take into account the kind of thing you are talking about.

You can use probabilistic events and take into account some of the uncertainties in the system, and then cause the system to adapt to those conditions, and then develop a new deployment schedule plan.

I have only talked about sealift. There is a whole airlift problem that is associated with this and incorporated into the problem. MSC's problem here is the sealift, but there is also an airlift problem. You have got the same kind of scheduling problem for airlift; most of the same rules apply. You also have the problem of making the connection between the two. A significant portion of cargo that can be lifted can go by air, and air has a very short response time. If you find you are falling behind, and you have a very critical situation, you can then redirect that airlift effort and use that to try to fill the gaps.

Interactive Vessel Scheduling

at EXXON

Thomas E. Baker

EXXON CCS
Florham Park, NJ

As we were talking earlier about optimization and how to deal with aggregation and with decomposition, I kept thinking of all of the things that we went through in developing our interactive scheduling system. And so on a case study basis, it may be interesting to the MSC people to see what we had to go through to solve this one particular problem. As we have said several times, this is one ship scheduling problem. And there are at least a dozen different varieties.

I will try to describe, for the purpose of the people who are concerned with modeling, which problem it is, but I think the flavor of the types of approximations that we make and the types of algorithms we try to use and how we try to interact with the user may be useful to the people who are looking at other problems.

Basically, this is an affiliate coastal supply problem (see Figure 6.1). The scheduler is sitting in a refinery. He has a dedicated fleet. He can load several products on each ship. He is to keep track of what is going on in the different terminals and keep them all within bounds. So, as opposed to just the window delivery type problem, he is really doing the inventory control for all of the terminals that are out there, and he has flexibility in terms of how much he drops off on each visit. He can drop off more on two visits to a terminal, or less in three visits. He has a frequency/drop type of flexibility, which is a very important part of the problem and which also makes it difficult.

INTERACTIVE VESSEL SCHEDULING

TYPICAL PROBLEM DIMENSIONS

COASTAL PRODUCT SUPPLY

- 1 SOURCE
- 4 VESSELS (DIFFERENT CHARACTERISTICS)
- 2 CHARTERS
- 23 TERMINALS
- 9 PRODUCTS
- 3 MONTH TIME HORIZON

≈ 150 PRODUCT-TERMINAL COMBINATIONS

~ 30 VESSEL ARRIVALS PER MONTH

FIGURE 6.1

The size of the problem we are talking about is: one source, four vessels with a couple of charters that he can bring in, 23 terminals, and nine products, although there are roughly only 150 product terminal combinations with which he is concerned. He used to work with a set of work sheets that were quite thick, where he did inventory projections for all of the different products in the different terminals and figured out where it was in the ships. He has thrown away his work sheets, which is good. The interactive system has been in production for the past year and has been deemed a big success, and now they are spreading it to three other refineries.

Every time in the past at EXXON when we tried to do this type of thing it failed, so people were watching this with great interest. There is a three-month time horizon, and maybe we will talk about that later. He also uses this system with a one-year time horizon to do facilities planning, but the basic operating schedule that he uses on a day-to-day basis runs for about a month to six weeks.

To give you an idea of the incidence of arrivals, in this model there are roughly 30 vessel arrivals at a terminal per month. So I guess the average leg of a voyage in this model is about four days.

See Figure 6.2. What we are dealing with in this problem is deterministic demand and single-stage distribution. We are looking at multi-port voyages, the variable frequency/quantity that I mentioned before, and no product-compartment assignments. We felt that we would not work on the problem if we needed product-compartment matching at this level. And it is actually not a constraint because they have so many compartments on these coastal products ships that they can run up and down and load things to suit when it comes time.

There is no vessel speed adjustment, although it can be put in as input. The scheduler is using the system for an operating schedule with a rolling time horizon. He is using it as a "what if" box and takes care of the problems as they occur.

INTERACTIVE VESSEL SCHEDULING PROBLEM CHARACTERISTICS

- DETERMINISTIC DEMANDS
- SINGLE STAGE DISTRIBUTION
- MULTI-PORT VOYAGES
- VARIABLE FREQUENCY/QUANTITY FOR VISITS/UNLOADING VOLUMES
- NO PRODUCT-COMPARTMENT MATCHING
- NO VESSEL SPEED ADJUSTMENT
- OPERATING SCHEDULE WITH ROLLING TIME HORIZON

FIGURE 6.2

TEMMLER: Would you elaborate on the assumption of single stage distribution?

BAKER: We are not concerned with the distribution problem beyond the terminals. We deliver to the coastal terminals and then it gets distributed, but we are not concerned with the demands beyond the terminals.

This system is part of an R&D project. Just a little background here (see Figure 6.3). It is called the Alcohol Project. We spend a lot of time on acronyms.

We are looking at more than just ship scheduling. We are looking at combinatorial scheduling problems in general.

Over the past four years we have done a number of prototype systems (see Figure 6.4). You also see the production systems we have in place. MIST is the one I am talking about now. There is another one we are trying to get going, which is also a ship scheduling one that is more like an assignment problem.

We have done a lot of work in process scheduling. We have six of these out in the field. We also have a helicopter dispatching problem, which is in place out in Australia.

All of these systems have the same basic philosophy, and we have been saying "We did it this way" all during this meeting. Now you get the whole full-blown description of how we did it. The basic philosophy of the system is that the human controls everything (see Figure 6.5). He has a deterministic simulator, or some calculation tool, which gives him his inventory projection and evaluation of the schedule.

By filling in the data he makes the simulator respond the way he wants it to, and we assume that this is the best model in the computer of the problem at hand.

Now, there are many other models. As you will see, there are LP models, network models, and sequencing models in here. They may deal with different parts of the problem. They generally check back

INTERACTIVE/ALGORITHMIC COMPUTING

The ALCOHOL Project

(ALGORITHMIC COMPUTING WITH HEURISTICS ON-LINE)

What is required to allow an experienced human scheduler to interact completely with an optimizing algorithm in the solution of combinatorial scheduling problems?

FIGURE 6.3

INTERACTIVE/ALGORITHMIC COMPUTING

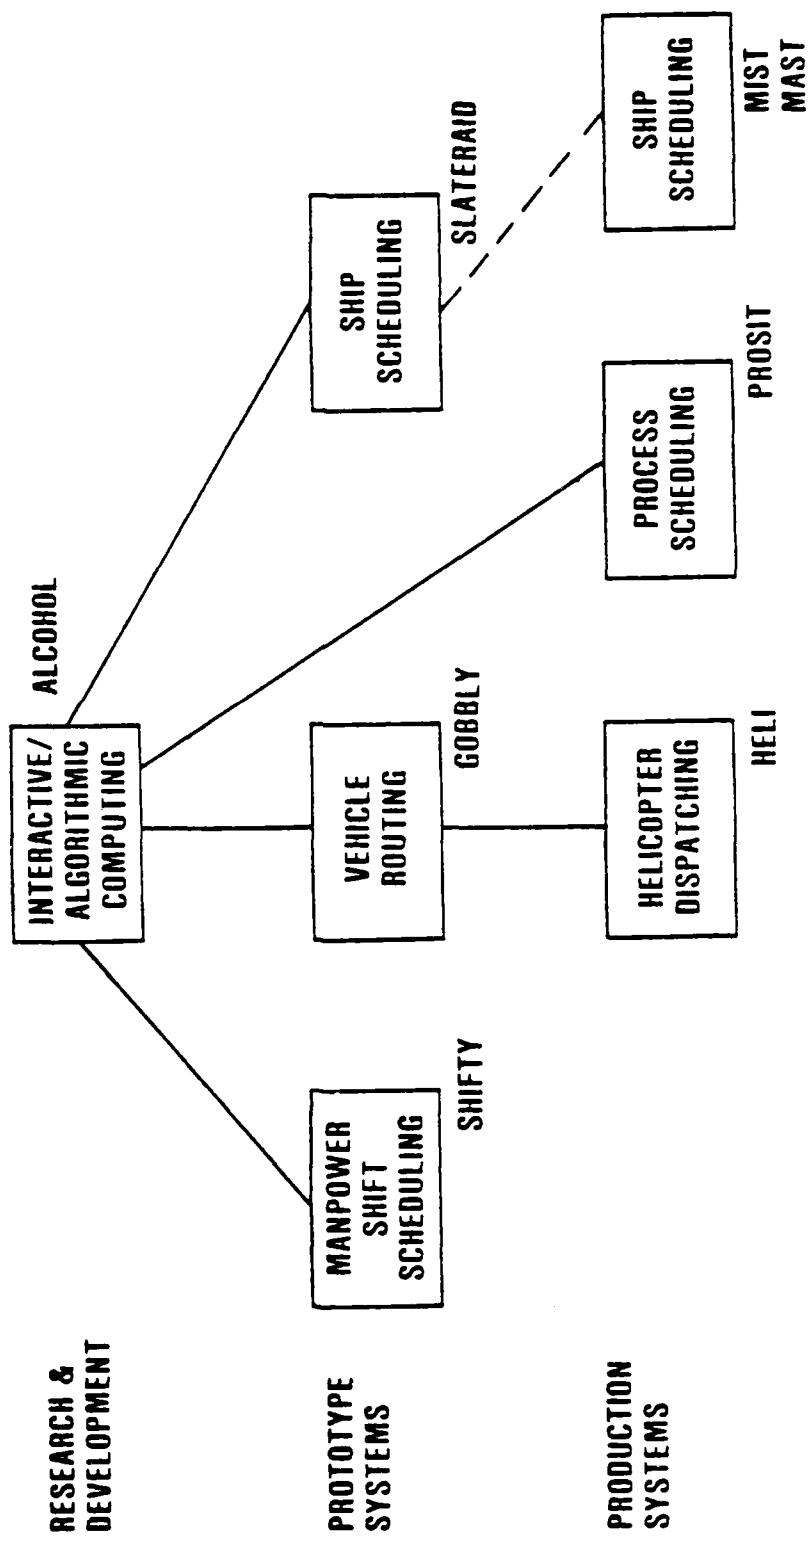


FIGURE 6.4

ALCOHOL PROJECT
(Algorithmic Computing with Heuristics On-Line)

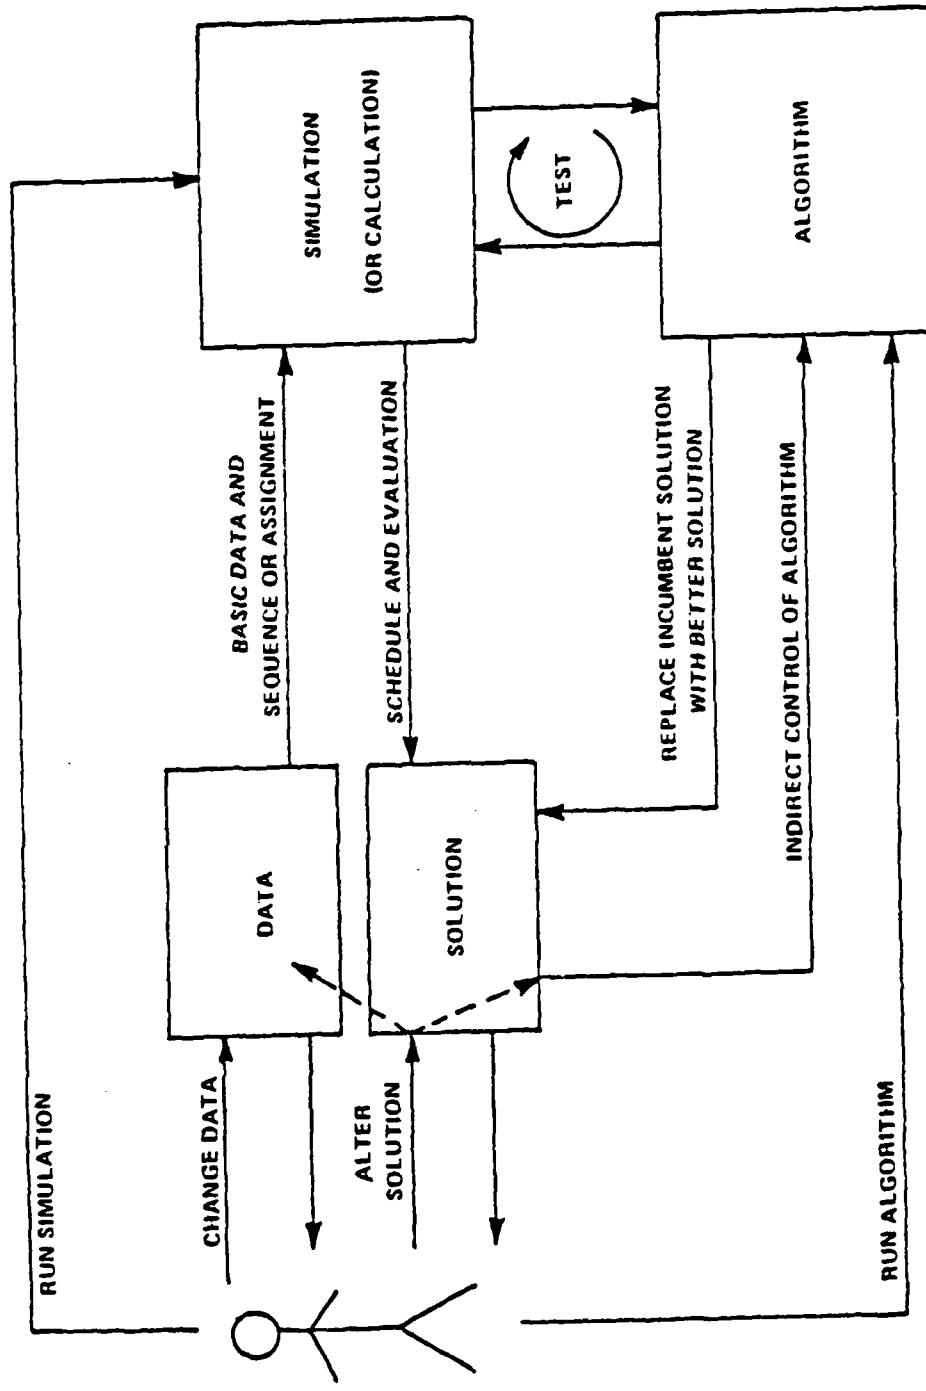


FIGURE 6.5

against this simulation model for a function evaluation to see how well they are doing.

Any heuristic that you want to use down at this level can always see how well it is doing by comparing its trial solution against the incumbent solution which is maintained.

In our system the heuristics, by definition, do not do any worse than the human. They can only do better. They will not replace him, and they will not take a step that does not improve what the human has already decided upon. I will talk about that later.

That is the basic flaw of the system. It is not really a model-based system in the sense that I have got a mixed integer programming model, which is the basis for the system, and with the human interacting with it. He is really interacting with this model up here, which is more like a simulation model.

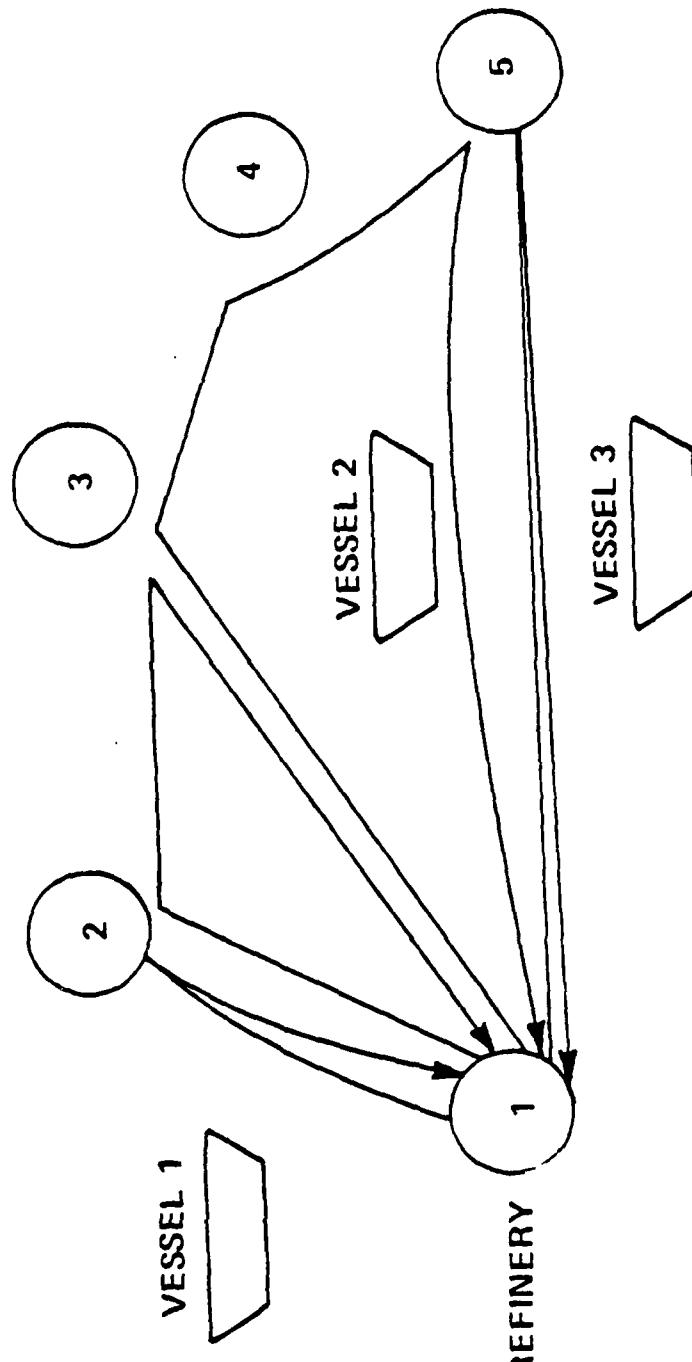
There are other mathematical models of the problem which are interacting. But he does not know that these things are going on; he just says do a loading or do a routing, etc. Those things are really transparent to him.

We took a look at the ship scheduling problems, and we tried to classify them into different types of problems that the scheduler has. One of them is a longer term problem of voyage selection and ship utilization (see Figure 6.6): do I bring in the charters; do I take the bigger ships and put them on the longer voyages; do I use them on the milk runs and use the small ships to fill the gaps; how do I get a balanced set of voyages?

That is one problem. The solution to that is LP which can be called by typing 'LP' on the command line. This was one of the first LPs ever formulated, I think (see Figure 6.7).

All of you have seen this formulation in your textbooks. Basically, we are assuming round-trip voyages. For example, column one shows vessel 1 going from the refinery, stopping at ports 2 and 3, and then coming back. The time of that voyage is here (8.2).

PROBLEM: VOYAGE SELECTION, VESSEL UTILIZATION



SOLUTION: LINEAR PROGRAMMING FORMULATION
'LP' ALGORITHM

FIGURE 6.6

LP FORMULATION

TERMINAL VOLUME BALANCE ROWS	VESSEL 1			VESSEL 2			VESSEL 3			RHS	
	1	2	3	4	5	6	7	8	9		
P1	13	13	13		8	8	8	8	19	\geq	
P2	6	13			8	4			8	\geq	
P3	7		13		4	8			8	\geq	
	
	NUMBER OF VISITS PER TERMINAL			MIN. NO. VISITS P1			TIME AVAIL. V1			NET DEMAND P1	
	P1			1 1 1			.7 10. 4.2			P2	
	P2			1 1 1			.7 10. 4.2			" V2	
	P3			1 1 1			.7 10. 4.2			" V2	

	VOYAGE TIME PER VESSEL MAXIMUM			TIME USED V1			MAX. NO. VOYAGES V1			NET DEMAND P2	
	V1			8.2 6.0 4.1			.7 10. 4.2			V2	
	V2			.			.			" V2	
	VOYAGE TIME PER VESSEL MINIMUM			.			.			MAX. NO. VOYAGES V1	
	V1			1 1 1			1 1 1			V2	
	V2			.			.			" V2	

FIGURE 6.7

We have the single loadings indicated in the last row. The scheduler has a voyages set; he has generated the possible voyages that he wants to consider. In this case it is not so bad. There are about 40 round trip voyages per vessel that he generally considers as a complete set.

So, for four vessels this becomes a small LP with less than 200 columns in it. There are voyage times for each vessel. There are other kinds of control rows, like the number of visits per terminal, which will be used later in the algorithms.

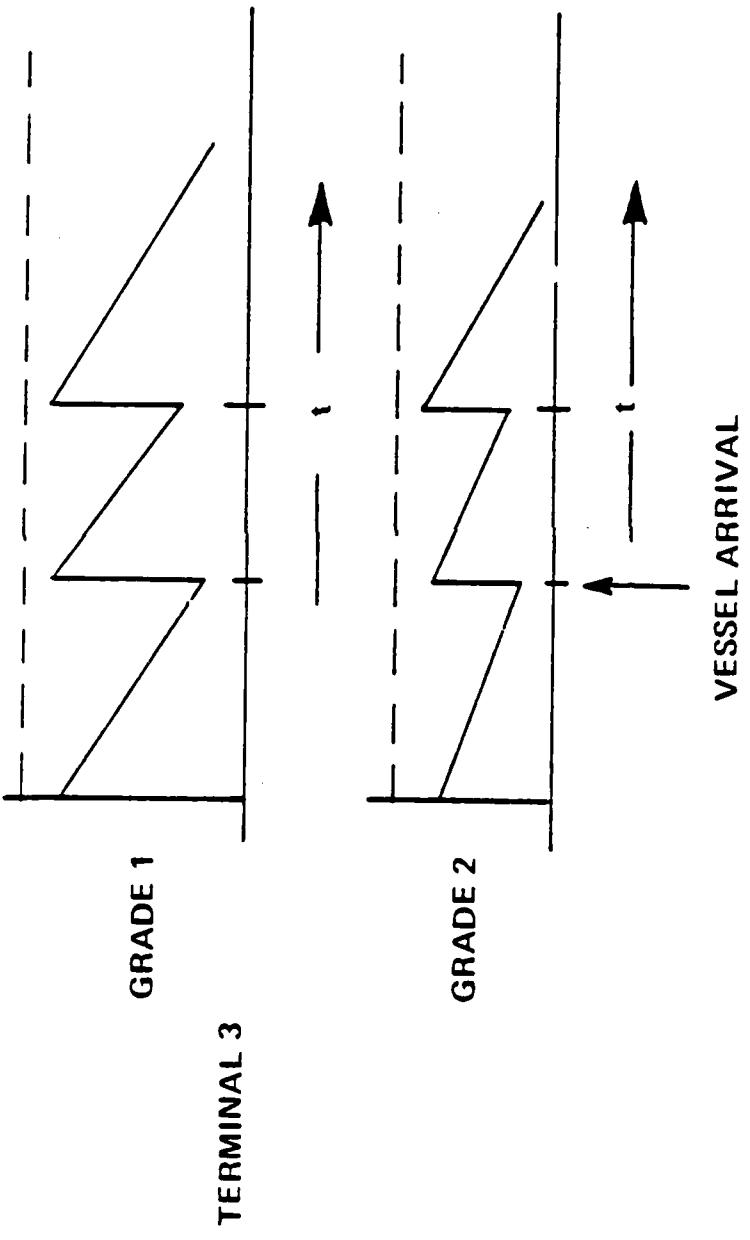
This is the LP that is generated. Later we will come back to this and adjust things to generate different types of voyage sets. But this is the basic model which was set up. We found that the user can interact with the LP model very easily, and I will talk about that more.

Suppose you have decided on a routing sequence for each vessel. Once you have done that, you basically have defined a distribution problem, given the capacity of the vessels, and that we model as a network in a multi-time period sense. What we are trying to do is follow, for each terminal, the inventories over time, and to make sure they are within bounds (see Figure 6.8). This takes a routing sequence, generates the distribution model, solves it, and gives the solution back in terms of the loadings for each vessel and the drops for each terminal.

The network model looks like that in Figure 6.9. Using network models for this type of interactive computing is very nice, because the answers come back so fast that the user does not even know that they are out there doing anything. This problem, for the nine products, has 1000 nodes and 2000 arcs, and it gets generated, solved -- in fact, it goes through some Lagrangian relaxation and has the solutions sent back, and it takes about the same amount of time as most of the data retrieval commands.

The basic model looks like this. There is a series of nodes for the vessel stops, and he can unload at terminal 2 because vessel 1 visits terminal 2, and then it goes on to the next terminals. This

PROBLEM: DETERMINE OPTIMAL VESSEL LOADING/UNLOADING FOR GIVEN SCHEDULE OF VISITS



SOLUTION: NETWORK FORMULATION WHICH CHECKS TERMINAL INVENTORIES BEFORE AND AFTER VESSEL ARRIVALS.

'LOAD' ALGORITHM – TOTAL VOLUME BASIS

'PLOAD' ALGORITHM – INDIVIDUAL GRADE BASIS

FIGURE 6.8

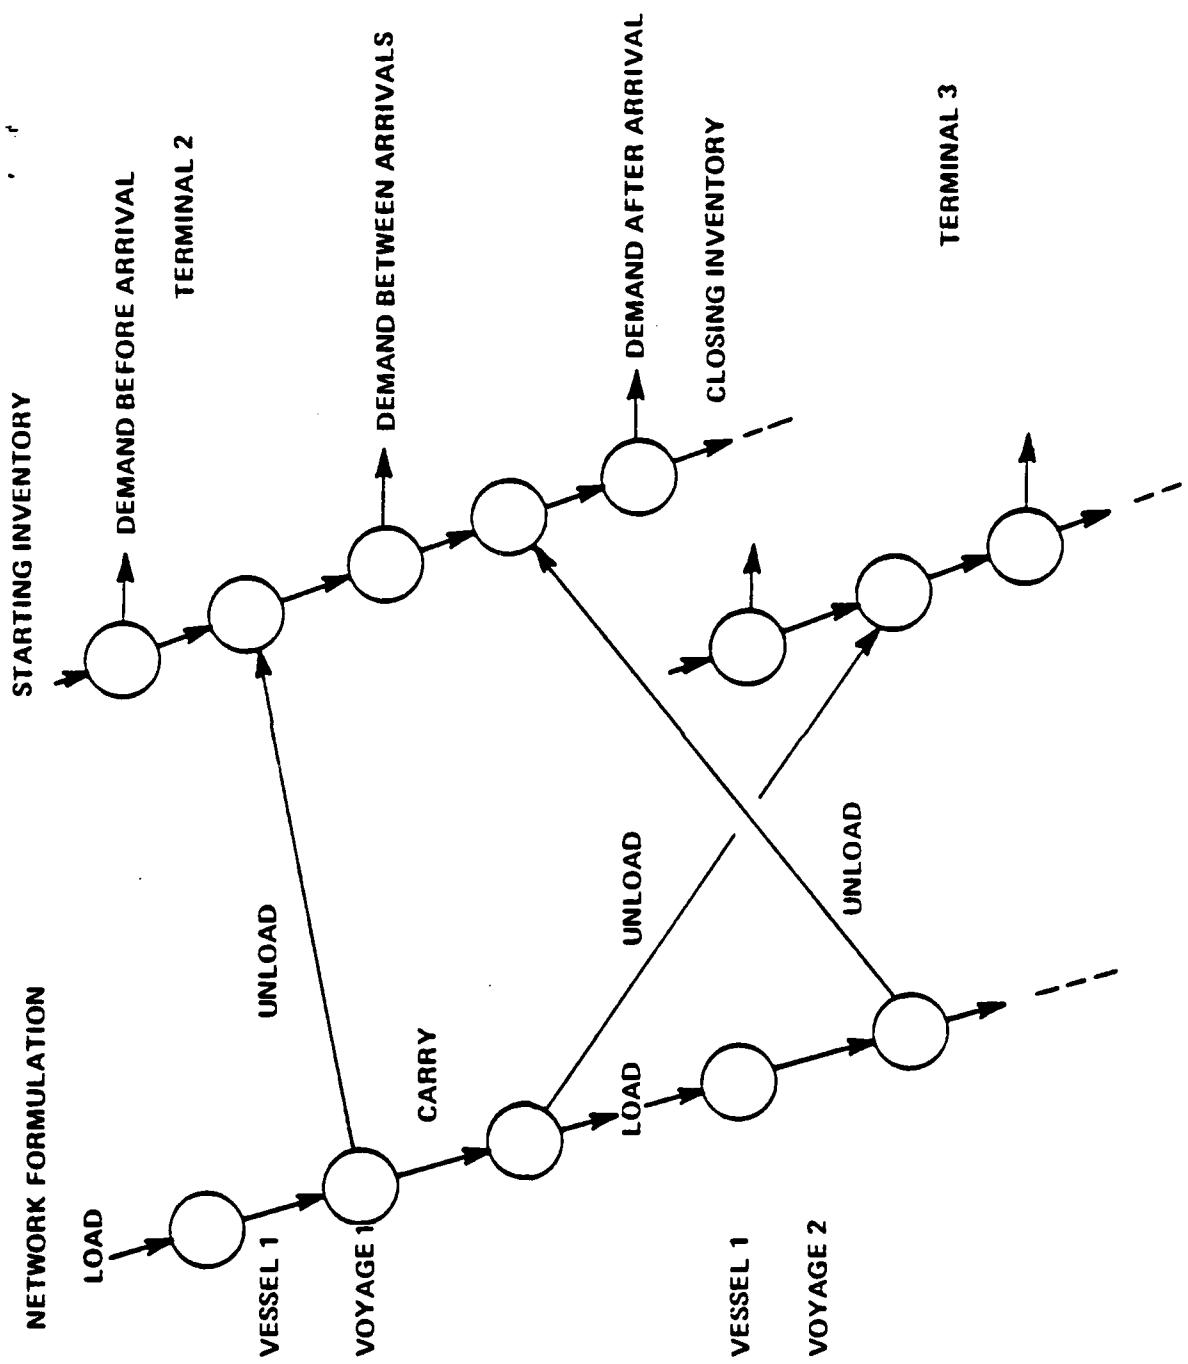


FIGURE 6.9

was a two-port voyage for vessel 1. You have a series of time periods, inventory nodes for the different products and the different terminals, and you check the inventory before and after the unloading, so that the whole program decides how much to unload for each voyage, and how much to load up.

The Lagrangian part deals with some non-network constraints. For example, if I have loaded dirty fuel on one voyage, I cannot load gasoline in the same compartment until I put some diesel fuel in there. Such constraints are non-network constraints.

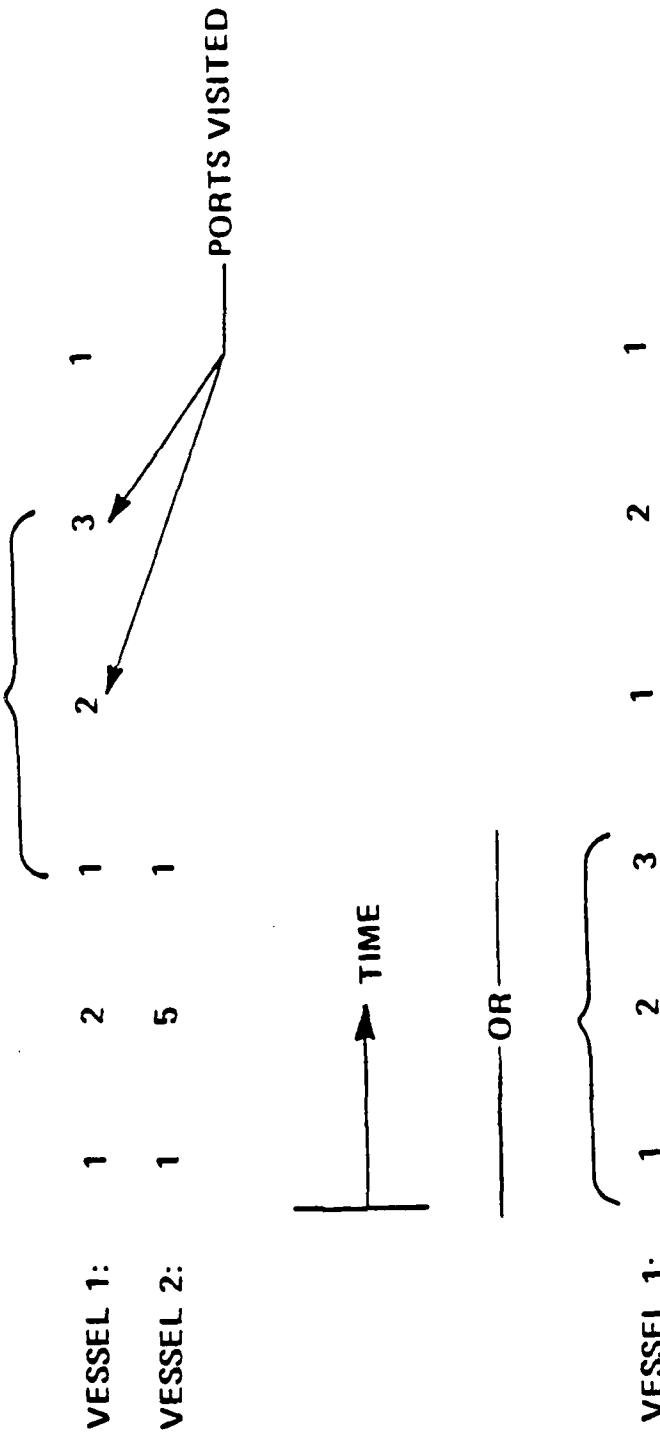
Another part of the problem is the sequencing (see Figure 6.10). This is not the sequencing of visits within a voyage; this is sequencing of voyages because I am choosing only round-trip voyages. Once I have done that, out of my LP I need to know whether I can run this round trip before the next round trip, and what that will do to my feasibility. The LP is very aggregated in time. This routing algorithm takes the round trip voyage set, treats it as a traveling salesman problem and uses Lin's method, re-sequences, and calls in loading, and so forth, so there is an embedded sequencing algorithm.

DOUGLAS: How does the size of the cargo get entered?

BAKER: Back in the LP, we assumed cargo sizes. After that, we always determine the cargo size through some network loading. So that was just an approximation to get a voyage set. As we worked down into the problem, we would keep adding more detail to it. If you tried to take all of the detail on the different cargo loadings and everything and put it up in that top problem, you would have a problem that would be too large for the computer.

I guess that will be clear from Figure 6.11. There is really an aggregation hierarchy. We decided not to aggregate geographically. We keep our 23 terminals or ports all the way through the problem. We toyed with that idea, and decided against it, but we aggregate everything else. We aggregate time. In the LP, we are considering one time period. We have lumped everything together by month, and

PROBLEM: SEQUENCE VOYAGES TO KEEP INVENTORIES FEASIBLE



SOLUTION:

TRAVELLING SALESMAN ALGORITHM (LIN'S METHOD)
WITH IMBEDDED SIMULATOR AND PRODUCT LOADING ALGORITHM

'ROUTE' ALGORITHM

FIGURE 6.10

INTERACTIVE VESSEL SCHEDULING AGGREGATION HIERARCHY

<u>ALGORITHM</u>	<u>TIME</u>	<u>LEVEL OF DETAIL</u>	<u>PRODUCT</u>
LP – VOYAGE SELECTION	BY MONTH	BY TOTAL VOLUME	
ROUTE – VOYAGE SEQUENCING	BY VOYAGE	BY TOTAL VOLUME	
LOAD – VESSEL LOADING INVENTORY CONTROL	BY VOYAGE	BY TOTAL VOLUME	
SIMX – SIMULATION COST EVALUATION	BY EVENT	BY TOTAL OR BY GRADE	
PLOAD – VESSEL LOADING INVENTORY CONTROL	BY EVENT	BY GRADE	

FIGURE 6.11

we break it down, then, by event, as we move on through the different algorithms. At the upper levels, we aggregate by product. In the beginning we do not consider individual products; we consider total volume. We select voyages on that basis, and we get down into the sequence on loading and add in more product, and then we start worrying about what network constraints of grade-to-grade switches we need and those kinds of things. We do not look at them at the upper levels.

So you can see, as we go down through the routing algorithm, the loading algorithms, and then finally down to the multiple-product loading algorithm we end up getting down to the real nuts and bolts, which are event by event, grade by grade, as to what is happening.

We have not aggregated resources such as ships. But I would guess in the MSC overall planning problem you would have to do that too. There are probably five dimensions, considering all of those, that you would probably want to take a look at aggregating for different parts of the problem.

Consider the problem stated in Figure 6.12. The user can run the LP by itself, and, say, do more visits for some terminal, and get another voyage set. As it turns out, he interacts with that kind of thing very well. We are amazed at how quickly he picks up those kinds of algorithms on an input-output basis--or he can enter "OPT" as we mentioned before, which really runs through the whole thing and tries to do it all (see Figure 6.13).

OPT will generate an LP, get the LP solution, try to make a feasible schedule out of routing the ships, using the routing algorithm, and that, in turn, calls the product loading on a total volume basis. If the system finds it has troubles because it is not visiting a terminal often enough, it will go back up here in subsystem 'LP,' add some to the right-hand side of the LP, get another voyage set, and then go down through again--this is all done automatically by 'ROUTE'--and then finally end up doing the multiproduct loading. The user can control this process through any state that he likes.

PROBLEM:

- GENERATE A PLANNING BASIS SCHEDULE FROM SCRATCH

SOLUTION:

'OPT' ALGORITHM

- RUN LP TO SELECT VOYAGES
- ANALYZE SOLUTION, ALTER RHS'S IN LP
- RERUN LP
- RUN ROUTE (WHICH RUNS LOAD AND SIMX)
- ANALYZE SCHEDULE, ALTER RHS'S IN LP
- RERUN LP
- RERUN ROUTE
- RUN PLOAD TO LOAD ALL GRADES

FIGURE 6.12

ALGORITHMS

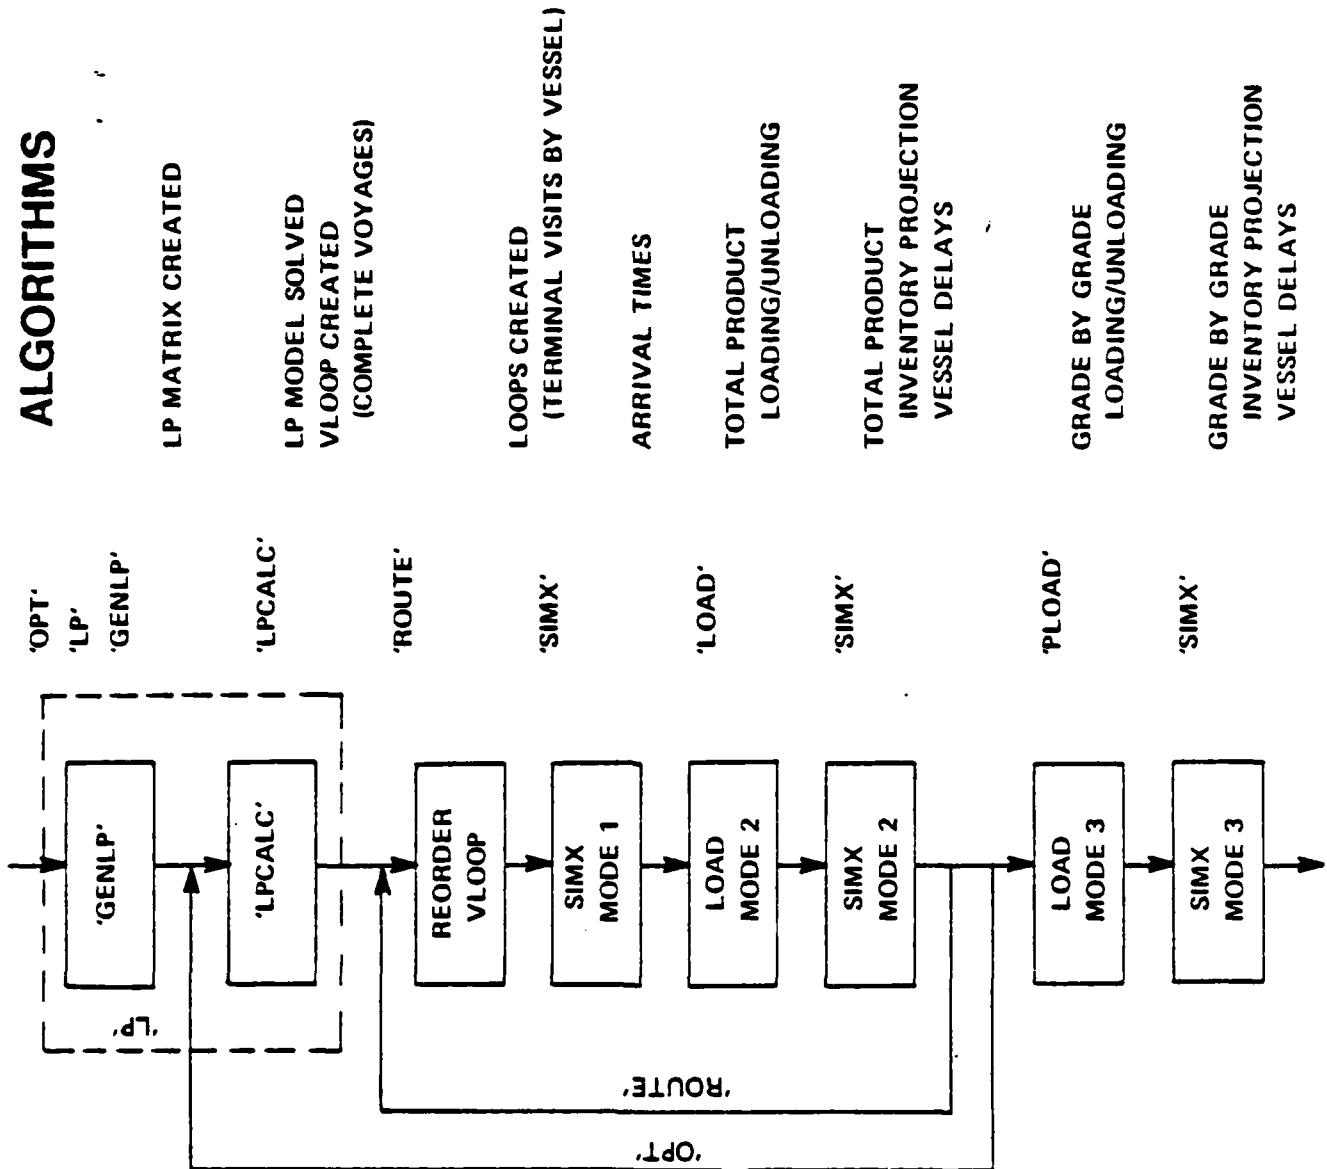


FIGURE 6.13

I will stop with Figure 6.14. I think everybody who deals with computerized routing problems says they save 15 percent of what it costs manually. I think that is an accepted number, and we also had a 15 percent reduction of operating costs. There are probably more benefits for us, though, in the ability for the user to do quick rescheduling without fouling himself up next month, and for him to do risk aversion, because he can now go through and see if he is running into safety stock levels at different points in time, and see what that means, and make interactive changes on that basis. The benefits, then, in here far outweigh the benefits in just having lower cost schedules.

The cost of developing the system was \$250,000, but we started with the ALCOHOL framework as an R&D project that has been going on for four years.

DISCUSSION

ENGLISH: Can you say anything about how this program helps the operator with respect to on-hire and off-hire for charters?

BAKER: For each of the vessels the user brings in a series of operating costs. The charters have higher operating costs associated with them, and the charter voyage sets are stuck into the LP as well, so the LP will not select those voyages, those voyage sets, if it can avoid them. Just naturally, solving that first LP will tell him whether he needs to bring in any charter vessels.

ENGLISH: Are these voyage charters?

BAKER: Yes. Each charter possibility is treated just like it is another vessel that he has in his fleet. The possible voyage sets for the charters are also in there, but they are in there at higher costs, so the LP will not select them unless it really has to have them, and that is the first thing the user normally looks at as he goes into the next quarter. He says, "Do I need any charters?" and he runs the LP, and it gives him a set of voyages. He sees

INTERACTIVE VESSEL SCHEDULING BENEFITS

- 15% – 20% REDUCTION OF OPERATING COSTS ON THE BASIS OF PARALLEL SCHEDULES
- UNKNOWN CREDITS FOR RISK AVERSION AND NEAR-OPTIMAL RESCHEDULING DUE TO UPSETS

FIGURE 6.14

charters in there. And then he knows he will not make it.

DOUGLAS: If he comes out with too many, i.e., what would be considered as too many charters, then might he take on a ship to cover some of that?

BAKER: That's right.

KASKIN: Are you working on this as a multiple source problem, or will there be interaction between the sources?

BAKER: Obviously, the LP formulation is what it is based on, and being able to reroute things, it treats round trip voyages. When you take that away, then you have a more difficult problem. In the planned extension of the system we have looked at doing multiple-sourcing. It is not so easy, because once you open up the traveling salesman problem for a ship not coming back to the same source, you have opened up a bigger problem.

General Discussion Following
the Invited Presentations of February 3

SOLAND: We are going to spend the next half hour in somewhat informal discussion with our speakers of today as a panel, with the one change that Walt Maclean of MARAD is substituting for Russ Stryker.

I see the goal of this particular panel discussion today as to try to assess where we are in terms of applying OR and computer modeling to problems in cargo ship routing and scheduling. Maybe we could start out by trying to see what we are doing well in the area of cargo ship routing and scheduling. Maybe it is not through modeling or computers, but maybe just seat of the pants. But are there some things that we are doing well? Why are we doing them well? What techniques are helping us to do them well?

Let us try to spend a few minutes on this. And, of course, the answers can come from the military side, the contingency planning side, or the commercial shipping side. So, we will just let anybody who wants to say something start the discussion.

RONEN: I will not say what we are doing well. I will say what I think we can do well. Based on the literature review that I presented, I think we can do well what I call long-term planning and probably medium-term planning, medium-term scheduling, again basically applying integer and mixed integer programming methods.

The state of the art in that area is relatively well developed, and we can do that well if there is sufficient cooperation from the users.

SOLAND: Jai Jaikumar has just talked about a model that I would characterize as very short-term.

RONEN: Do you want me to answer that?

SOLAND: Maybe it is just a hope that we can also do well on short-term planning.

RONEN: That model was more applicable to vehicle scheduling than to ship scheduling. That is the major problem as far as I see it. He tried to compare it to the model of Appelgren, which treats a medium-term scheduling problem with ships. He compared it to medium-term scheduling models in shipping, not short-term scheduling problems.

BISHOP: I do not really want to challenge Professor Ronen, but I would like to make a comment on long-term planning, at least in the tanker business. When you start talking long term, even as he defines it, we do not know what our requirements are. And that is really a problem. You do not know what you have to move. We have a plan of what we have to move. But somebody will say "If you try to schedule out that far the schedule is no good, because what you had to move is not what it was anymore."

RONEN: That is the short-term problem, what to do once you get to the point when you have to execute the plan. You can plan ahead. The question is how are you going to carry out the plan in a changing environment.

BISHOP: What I am saying is that when you try to plan it out you really do not gain anything. You went to all of that effort to try to figure out what you are going to do three months from now, and three months from now you do not have the same problem.

SOLAND: That is the uncertainty coming in?

BISHOP: Yes. Things have changed too much. It may be a characteristic of our business, but it is certainly a problem we perpetually have.

TEMMLER: You asked what are some of the good things we might be doing as brought out in some of the presentations. I think I can see some that can help me, certainly, and probably others.

One is, as Carroll Keyfauver brought out, to reduce the size of the problem. You then get a sequence of much smaller problems. It is the old axiom that someone has in his office everywhere one of us works: "The problem, when reduced to its smallest," etc., "can be controlled."

And as Walt Maclean has said, at least once already, we have to define the problem. Maybe that should have come first, before getting into it. I do not think we are good at defining the problems. I think bringing that out has been a good point today. Maybe the problem has to be defined better, in terms of the objective, whether it is to get a profit, whether it is to get a return on something, or whatever. Somebody has to make the statement.

And as Jai Jaikumar said in his presentation, they produced the scheduling system because the chairman wanted it. There is the point that if the chairman wants something, he gets it, whether there is a good objective or not.

I think those good things can be seen as long as you realize that there is the other side of the coin, that maybe the problem has not been defined well in some of these situations.

WEBSTER: I have a picture on my wall that says "For every complex problem there is a simple solution, and it is usually wrong." I think this is one of these cases, too. I would like to amplify on what you said.

If you deal, as I do to a certain extent, with a real shipping line with real problems, I think you have to face up to the fact that in the foreseeable future you probably will not be able to quantify and to crystallize, in either terms of constraints or logic, what is required to meet the real world situation. It would be a mistake to aim, at least in the near term, for a programming method which does it all, which does not have the man in the loop, which does not have the ability for us to tap these people who have been dealing with these problems for a long time--to tap this experience and to incorporate that in some interactive mode.

We have had discussions here of methods in which there is some interaction. But the general trend or tendency, as I hear it, is towards solutions which are not interactive. And I would say that, based on my experience, that is unlikely to be very successful in the foreseeable future.

SOLAND: You brought up an interesting point, the use of interactive systems. And we have two people here in the audience, Tom Baker and Burnie Douglas, who have been intimately connected with interactive systems. And I think I will throw this out as another question. Is this or is this not the way to go to get substantial improvement? And how about the size of the data bases which are necessary for comprehensive systems of this sort?

CHRISTOFIDES: There is a lot of difference between shipping and vehicle scheduling. There the short term is very short term. It is something that very often cannot use the interactive mode. It is often the case that you will have a breakdown of a vehicle. The system has to have the facility for a human to make changes, but not in the decision-making loop.

SOLAND: Are you saying that there is a difference with ship scheduling?

CHRISTOFIDES: I suspect that in shifting the horizon, what is called the short term may be considerably longer than in vehicle routing.

SOLAND: Captain Scott talked about making decisions within an hour, or sometimes several hours. I think maybe that is a question worth further discussion.

MACLEAN: It seems to me we have talked about a wide variety of problems today, and so far we have not gotten any structure together in which those different problems, relating to the transport of goods by sea, appear in a matrix of some sort.

I think, depending upon who you are in the system, you look at the problem differently. The guy who owns the cargo is probably interested in traffic management in terms of trying to see how he can get his cargo onto available rail cars, trucks, ships, to get it where he wants to. He has one problem.

The guy who is down there operating the ship has another problem. He has a perceived schedule commitment that he is trying to meet, and he does not really see what kind of service he is trying to provide, except within the constraints that have been set out for him.

We have talked about the idea of simplifying the problem. And in actual practice, that is what we have done.

We do not have any comprehensive view. We do not have a view of what the market analysis says, what the corporate strategy is going to be, how much market share they decide they are going to try to get on a trade route, what types of assets they are going to be able to employ, what kind of return they have to get, what is the procedure by which they are going to formulate a service, and how do they then propose to carry out the structure and the function to make it happen.

Until we start putting this into a framework, I think to try to formulate it is going to be very difficult for us. I think all of those things have to come into play, whether we are going to talk about an interactive system or somebody sitting off with his computer spitting out recommendations periodically for somebody to buy. I think that is what we have to come to grips with.

The Navy problem I think is quite different from the commercial problem, although there are clearly interfaces where they come together with common purpose.

Take this question of minimizing the cost. It may very well be that that is a primary function because of the limit on budgets. But on the other hand, it is really a matter of trying to maximize the utilization of the assets you have, however you want to define those assets, for the purposes that you are charged.

It seems to me that we have got to build this framework before we can really attack it in terms of formulating solutions.

FOARD: I think the military problem encompasses it all. There are basically three different problems the military is talking about today. One is the long-term planning function, evaluating the capability to deliver a bunch of cargo from here to there over time, i.e., the planning function.

And then we get into the execution function, which is where Jai's presentation came in. We have a requirement to take three or four different courses of action, and then we can assess, in a rapid turnaround manner, the gross feasibility of the courses of action. That's the second part of it.

Now, the third part of it is when the decision is made to execute one of those courses of action. We only have to get into the short-term scheduling on a day-to-day basis, e.g., given the cargo is made available on a given day and a ship breaks down the next day, who's going to pick up that cargo. We have each one of those three problems here. In different ways, they can be related to each one of the commercial problems in the same way.

SOLAND: Addressing Walt Maclean's points, perhaps it is a question of stipulating objectives and criteria, not necessarily one. Maybe this is the first step that is needed. Once the criteria to be considered are well understood and established, anyone can try to choose either one as an objective function and the others as constraints, or to go into a multiple criteria decision making model.

Some of what I heard, in terms of meeting required delivery dates, brought up the subject in my mind of goal programming, paying penalties for being late, and increased penalties for being later. Maybe that is one way of handling one of the criteria there, which is getting what you need to where it has to go, and on time.

KASKIN: Well, I think it is apparent that we have to have some definite objective functions defined. And I think that during the discussion sessions tomorrow we will present, at least in the military side, possible objective functions. And then we can interact and maybe get some better definitions, if they are still too fuzzy, in order to see exactly in what direction we are trying to go.

T. BAKER: I am a little amazed that we might be even looking for a philosopher's stone. Within our company, I can think of at least four different types of ship scheduling problems that run the gamut from a multiple traveling salesman type problem to an inventory management problem to a more or less straight type of assignment problem, and all kinds of variations in between.

I think if you want to try to classify them, you can go back and look at all of the OR literature and talk about, e.g., job shop problems with due dates and multiple traveling salesmen problems, and you have the whole gamut if you are going to take the whole area of ship scheduling. And I think you are going to be about as successful at coming up with one technique that will solve all those problems as anybody else has been.

SOLAND: But at the same time, you seem to be suggesting that there are a number of problems which can be cast into somewhat familiar molds of mathematical models, maybe with additional complications and so forth.

T. BAKER: That is true.

MACLEAN: What you are suggesting is that we really need to identify what the driving forces are, and out of these driving forces we can see a mode of operation develop that leads to some objectives. And the objectives probably will be different for each one of the driving forces. But maybe the mechanism for going through it will be parallel in some respects.

T. BAKER: Yes. I guess our approach to common scheduling problems has been different. They are not model-based systems in the sense that there is a single model, like the one we just heard about, underlying the whole thing that you are working with. They may have a whole collection of models. And you use a model for the aspect of the problem which it fits best and work with them in that fashion.

And there is the recognition that there are so many different characteristics to these problems that you cannot, except in special circumstances, grab hold of the whole thing. It is just not possible.

MACLEAN: You suggest that you need to have an executive program with a whole series of options that call up a bunch of subroutines, depending upon the nature of the particular problem.

T. BAKER: I know academics do not like to think about that kind of thing. But yes, basically, that is it.

SOLAND: Is that the way your system is structured?

T. BAKER: Yes

VOICE: Isn't that the human part? Aren't you describing an interactive system?

T. BAKER: The human is in there, too. But even in a ship scheduling system, where he would say just "optimize," i.e., he wants to run the whole thing from scratch, it is not one algorithm, but a series of algorithms. We just gave up on the idea of trying to solve the thing at once.

SOLAND: Do you think that a somewhat similar approach might work for the contingency planning problem or the strategic deployment problem?

T. BAKER: Yes, I do.

SOLAND: I think the MSC people will want to talk to you.

T. BAKER: I am having trouble seeing how they are going to do it any other way.

HALL: I would not want to speak for the MSC people. I am from the Joint Deployment Agency, and we have a broader problem. Sealift is part of it. We also have to worry about integrating the airlift with sealift, as well as the intra-CONUS moves and the intra-theatre moves. We certainly are not looking for one model or one algorithm that is going to do all that. I do not think anybody intended to give that impression.

I think the question really is what are the applicable techniques, the ones that have been successfully used to solve the pieces of the problem. Maybe we have not properly stated what we think our question is. I do not know. I would like to hear somebody from MSC address that. I do not think anyone is really looking for one model that is going to do the whole thing.

JAIKUMAR: I would like to address the question which Tom Baker raised. Suppose you have a number of small models, and I agree with your philosophy that you need a number of small models, even though the model I am talking about was fairly integrated. But the question I ask is this: is it ten small models or 50 small models? What degree of aggregation are you going to deal with? How many different decision variables are going to come into one model?

You cannot have either ten, 20, or 50 small models. What you are trying to do, really, is try to see when these interfaces can be integrated, how does one affect the other. And hopefully, we could, as we go along, start stitching some of these models together so that you could, in fact, get integrated. You are never going to get to one model.

But one of the things which I am interested in is to really see how these models get stitched together, how they get integrated, which itself is an interesting and quite practical question. In the Air Products situation, we are looking for control in transportation, and we have done it through this model. But, in fact, different parts could have been separate.

So, the question is not in terms of philosophy, whether you have a number of small models or several big ones, but where you draw the lines when you cut up the problem.

SOLAND: That is definitely a good question. We can perhaps say that every optimization is a suboptimization because we have not gone to the biggest system yet. But we cannot work any other way, can we?

JAIKUMAR: I think it is a building block type of thing. You start off with small models, but eventually you try to put them together, and you try to see ways in which you can, in fact, do that.

T. BAKER: I will give an example. You take a supply system where you have supply and demand on the ends, and your ships are doing the supplying. Even if you set up a full-blown, mixed-integer program formulation of the whole problem, which we have done in the past, depending on whether you are supply-constrained, demand-constrained, or ship-constrained, the model will operate in a very different fashion. And if you are using branch and bound on it, the branch and bound search, even though you have got the same number of coefficients, will be quite different. And you will want completely different strategies to try to solve it, depending on whether you are really ship-constrained or supply-constrained, or whatever the physical problem happens to be, even though it is the same model.

JAIKUMAR: Isn't that a valuable insight, that you do operate with different strategies? But the fact is that, looking at it as a large model, you do different things strategically. I think that itself is valuable, to know if you are demand-constrained you operate in this way, etc.

T. BAKER: But picture a client who has been running the model during the winter and goes into the spring season with a change in demands, and now he cannot get good solutions anymore. It is hard to explain to him why the model is running ten times longer than it used to.

What we learned from that is that maybe we want to think about different models because they are different problems. In one case, with one set of demands, it is really a routing problem because it is the routing constraints that are establishing the solution. In the other, there is a capacity constraint on one of the ends, and then it is more like an allocation problem, even though it is the same model.

What it led us to look at is completely different models, not one large model

DOUGLAS: If it is capacity-constrained, don't you add capacity and see what happens or see how much capacity you have to add?

T. BAKER: These are operational scheduling models. We are trying to move LPG from Australia, to keep it within bounds so we can find out where to put it in the world. We may be long on supply or short on supply, or long on shipping capacity or short on shipping capacity. And the LPG ships are not interchangeable. When you have them, you have them.

RONEN: And there are very few of them.

T. BAKER: From quarter to quarter, that same problem has completely different characteristics.

JAIKUMAR: Are you thinking of, in the future, somehow integrating some of these models? Are you going to retain them and fine tune what you have? What direction will you follow?

T. BAKER: If I were funding academic research for the future, and this week I am not, I would say the more of these problems you look at, the more of them that are of job shop types, the more they look the same. I would look for algorithms which you could use for that general class of problems and characterize better the different types of problems that you could operate on and use the algorithms as modules, as operators.

YOUNG: In the only two applications we have heard about, heuristic methods play a key role in getting the solution, in one case the only role, and in the other case in getting an upper bound which allows you to even consider an integer programming model. It seems like we are reduced to heuristics in terms of addressing these problems. Maybe other people have different experiences.

BISHOP: If you really are reduced to heuristics, which is what our schedulers are using now, at least on simple problems the oil companies have with respect to supply and distribution, it is probably not worth a lot to try using computers, because you can get people to do it. That may change if the problem is large. Right now, we have guys who have been doing it for years. They come up with reasonably good solutions.

COPPINS: I found the same thing over the summer when I was consulting. It was a shipping problem. I think the truck scheduling problem, in contrast, tends to have so much freedom that you really have to get in there. You really can optimize. But as someone said this morning--if you only have two or three ships, there is not a lot you can do.

The problem I see is connected with the word "optimization;" it's been bothering me. You addressed it, and nobody said anything about it, and so I bit my tongue. But I do not think there is an objective function. It is hard to define one for a private concern. I think it is going to be extremely difficult in the military. And if you do, it is going to be a surrogate objective function. The question is do you want to optimize that anyway?

Really, what you are trying to do is to deliver the stuff on time, where you need it. It is nice to talk about minimizing costs, at least during a peacetime setting. I find it difficult to swallow the fact that it will be anything other than a low-level consideration during a conflict. As a taxpayer, that would bother me. I assume that most people would feel the same way. It is a multi-criteria problem really.

Maybe what we are looking for is a solution with which we are satisfied. What is a good solution that achieves the deliveries that we want? Cost is one aspect of it. I think achieving the feasibility is a main part of it.

I keep hearing "optimization" over and over again. We can talk about it in the private sector perhaps, but I am not sure that is an appropriate thought pattern for the military. That may be causing problems.

BALLOU: As Captain Scott mentioned, periodically there is some consideration of moving equipment, forces, and people around. Maybe it is not a major consideration in the cost sense, but it is certainly a consideration that we do not want to spend the whole DoD budget on it in that two-week period.

And so, what I have seen in the short time I have been involved is that, increasingly in the play of the exercise, we are starting to get asked dollar questions. It seems to me that previously you never even heard about it. But now they are saying, "Well, okay. How much does it cost to charter a ship to do it?"--or ten ships, or this sort of thing.

COPPINS: You are talking about peacetime?

BALLOU: I am talking about what is required to do the job. I do not know what is peacetime or wartime.

COPPINS: It sounds like there are different purposes, though.

BALLOU: I think Woody Hall was talking about it in the various scenarios that we are considering. There are different purposes even in stepping through the ways to handle a crisis. As Captain Scott mentioned, there are different objectives at each one of those phases, and if you are down to the hard scheduling, that is a different idea than if you were just trying to say "do we need to bring on extra charters? Do we need to do this or that?" There are different kinds of information that you are reaching for.

GALLAGHER: I would like to address the optimization problem. I am in long-term planning with the 1,000 ships or the 10,000 or 15,000 units that we

move. That is correct. Military value does not come into it at all. The optimization along those lines is in terms of our assets and the capability of using the assets to make the RDDs. If you were to ask me, my cost becomes my delivery date and the use of the assets. But the objective is to utilize those assets at the optimum to meet the RDDs.

Someone before suggested assessing penalties instead of using money and value, so we would begin to assess penalty points for missing the RDD just like you would assess penalty points if you ran over your budget. I think it is very much applicable. The optimization is definitely something we need even in long term planning, whether it is assets or something else.

KASKIN: In the same vein, in many cases we may not have an infinite amount of resources to draw upon. We have different phases and different availability of ships. Some are easy to come by; some require major political decisions, requisitions, et cetera, if a national emergency is declared. Therefore, given certain levels of resources that are available, we want to use the fewest number of resources necessary in order to satisfy those required delivery dates.

There is a tendency to try to use the least resources in that problem. We do have a model now that will give you a feasible solution. Now, if it says that it is infeasible, it may still really be feasible, but it is just the model. It just shows that given the resources that it assumes you have, that you still cannot do it, although a better model may show that it is feasible.

COPPINS: There are a lot of different goals involved. There is not a function to be optimized. It is a vector value function and these are all facets of it. One of the things that I am hearing is that you want really to do some sensitivity analysis so you can plan. There is not an optimal solution. You really have to do contingency planning. That is some of the stuff I ran into over the summer. You do not know what you are doing over the long term. You need to toy with numbers and parameters and see what it seems we should be thinking about.

KASKIN: The deliberate planning process is a capability that we need. The execution is where the cargo is coming to the port and that is where the ships are. We have all of the problems to solve.

COPPINS: They are different problems. You may have to divorce them from one another. You have to stitch the models together appropriately.

SOLAND: Jai addressed that too. If I may address that for a minute, just in terms of short-term problems, we can suppose we have to do a scheduling job and it is involved in this and this and this. Should we try to use some kind of optimization model? Should we do it completely by the seat of our pants? Should we do it by a heuristic method that just has some good basis behind it, or should we use a heuristic that is based upon an optimization concept, which is partially along the lines of what Jai Jaikumar has said? So there are three or four different possibilities just in the concrete, short-term problem. These are some things, perhaps, to keep in mind for tomorrow.

MACLEAN: One of the questions that bothers me, however, when you start talking about carrying out a sensitivity analysis, is that it begs the issue as to how reliable the data you have to work with may be. That is a question I think that would be worthwhile talking about; do we know enough about the specifics of cargo movements and ship capabilities and the environment that we are working in that we would be able to exercise any kind of sensitivity analysis?

COPPINS: That is why you do the sensitivity analysis; that is precisely the purpose of it. You do not know these data and so if you rely on one final output, that leaves you open to variations in the data that you cannot control. That is why you want to play with this stuff. The stuff that you are not sure of you twiddle and say "okay, what happens if I make different assumptions? Am I out in left field once and out in right field once? Is there anything we can do?"

I think that is a major problem. You have got to be able to do sensitivity analysis, because there is so much that is uncertain. You cannot wholly rely on one set of data and assumptions.

WEBSTER: You asked a question before about objective functions, and I think what I said previously about interaction is essentially aimed at the same kind of thing. I think that in the private sector defining your objective function is a major point because there is no one in the shipping lines I am involved in who really knows what you are talking about when you say objective function. Then once you explain what it is to them and they understand the concept, they are unable to quantify it and you are unable to extract from them what it is that they really want because it is a complicated arrangement.

My feeling is that you always have that kind of difficulty in setting up an objective function. The most successful thing I have been able to do is write ranking matrices which are really penalty points. It has been very difficult to find out what it is that people want. Everybody agrees on the profit or return on investment, but when you have to express that in terms of routing and scheduling, how much profit does it cost you for not picking up somebody's cargo and you have alienated his customer. There is a whole series of interesting questions that are very, very difficult to answer. We can all say motherhood is fine, profit is good, but to quantify it in terms of a matrix is very, very difficult in the private sector.

MACLEAN: You give a marketing man a goal to provide so many containers of cargo at such and such a port to be picked up by such and such ship, he goes out and knocks on all of the doors and twists arms and gets everything on board, and all of a sudden you have an overshoot of a very important shipper a port ahead of you and you have got to shut out something. Whose cargo do you shut out and who gets disappointed, and what does this do to all of your planning?

WEBSTER: I doubt if you can quantify that. It is pol...ics. It is corporate policy. It has nothing to do with rationality.

SOLAND: I think we might agree that at a certain upper level things are rather fuzzy, sometimes difficult to quantify or to get agreement on. But when it comes down to putting out a schedule that the ships are going to follow, it has got to be done. And we still want to focus on different ways, either heuristic or optimization-based, or perhaps other ways, that you do it. People are doing it. People will have to do it.

CONWAY: I hope you won't mind if I say something a little bit different. We are a pure container operator. I have been struggling all day to grasp some of the concepts that you are dealing with. I would be much happier if the word "ship" were taken right out of this symposium. As a pure container operator, a ship is not that vital to me. A ship is an expense item; a container is a revenue item.

It is very easy for a pure container operator to set up his schedule. If you want to compete in the commercial field, you pretty much have to have a weekly sailing on whatever trade route you are running; therefore your schedule is in multiples of seven. If you have three ships, then every 21 days a ship has to be back to the port it started at. If you have four ships it has to be back in 28 days. If you are going to the Far East it has to be back in 70 or 77 days, or 63 days or whatever it is. That makes the rest of the schedule rather simple. You know exactly where the ship is going to be at any given time. Ship scheduling is simple. The thing that drives us is the marketing and the selling.

I have been struggling with the idea of what schedules really mean. Schedule contains a repetitive element. That is why I say that the schedule itself does not seem that important. It is automatic. It has to be there for us to do our thing.

I am not sure yet how we can help with this problem. It would seem to me that in an emergency the government needs to pay more attention to scheduling the cargo than the ships. Scheduling the container is important--the ship is no good whatsoever without the container. I think a lot of that kind of work can be done ahead of time. I think we should be scheduling the cargo flow rather than the vessel. The problem is to get large amounts of supplies where they are needed in the right sequence. We need to apply the computer to cargo schedules rather than ship schedules.

Discussion Group 1 - Industrial Operations

Leader: Walter M. Maclean

National Maritime Research Center
Maritime Administration
Kings Point, NY

MACLEAN: The topic we want to talk about in this group is industrial operations. I think that there are really three questions that we want to address, which in turn have a series of subquestions associated with them. The first question is: What is the state of the art? The second question is: What are the problems as we presently see them? And the third question is: What do we think can or should be done in the way of improving our capability to schedule, optimally or otherwise?

It seems to me that the first question regarding what is the state of the art requires that we identify the driving forces. For what purpose are we scheduling? Are we scheduling for the chairman of the board, who wants to get this process organized instead of being carried out in an ad hoc or in a manual fashion? Or is there a market survey and a corporate strategy that has been developed that has to be responded to in some systematic fashion? Or is there some other driving force behind ship scheduling? Could we start out by identifying that?

BISHOP: At Chevron the scheduling is on a sufficiently low level that it really does not make a lot of difference what the policy is that determined what needed to be scheduled. The fact is that there are some guys sitting down there who have got ships and cargoes to move. And why this cargo has to go in one direction really does not influence what they do. Their routine and their approach to the problem is identical. They are given this to move to here and that is the scheduling problem that we are looking at.

MACLEAN: So it really is the cargo operation that is the primary mover.

BISHOP: Oh, definitely.

MACLEAN: And what does the cargo operation respond to? Your refinery output? Your refinery input? The market that is served by the refinery?

BISHOP: It is directly driven by the refinery's requirements. The refineries call up and say they need so much crude. Now the reason they determine that they need so much crude comes from a higher level. Management, all the way up to the chairman of the board, decides how much crude we are going to run, what grades we are going to run, and the refineries have to shake it out. Then they can call up and we give them the ships to move it.

MACLEAN: So presumably, the refineries, then, are responding to a drive from a market analysis.

BISHOP: More from sales. They go out and they measure the content of the product tanks and find they are low. They are going to need more crude. To the extent that they are not moving their products, or to the extent they can buy products cheaper than they can buy crude and refine it, they do not ask for crude.

JAIKUMAR: What is the lead time from the time the refineries make a request for crude and the scheduler has to react to when they need it?

BISHOP: Initially, it is done about two months ahead of time. If you are operating a VLCC slow speed from the Persian Gulf around the cape, you have got a 45-day travel time. But the schedule does not get finalized until, say, essentially the same month. With the short-haul cruises like the Mexican and the U.S., lead times may be three or four days.

DOENGES: Are you just moving crude?

BISHOP: Our department moves crude. There are other people that have products to move. But sea-borne movement of products tends not to be a problem. You get locked into such weird situations of tankage, travel, and constraints, that there are not many scheduling options. Vessels tend to get dedicated to shuttling back and forth.

I would expect that is also the case in the military, or at least in a lot of it. If you are serving islands in the Pacific with tankers, dropping off little parcels here and there, you get locked in by tankage constraints and you struggle around and find out what is the optimum pattern and then you dedicate a ship to it. Regardless of how efficiently the ship is being used, it comes down to the fact that it is the only way you can do it.

MACLEAN: Do you ever get involved in decisions of sizing the fleet?

BISHOP: Yes, that is done three times a year. Basically, we examine the situation and plan out what we think our requirements are going to be. Then we take a look at the number of ships that we have to do what we decide to do. Basically, we use previous experience, and we run a lot of economic analyses.

We do not have a full-circle scheduling model. We have programs that will simulate a ship running and the trade and give you how much it costs. Essentially, it is done in my shop. I have two analysts who sit down with computer output and they decide what is the optimum way to schedule the fleet. And then that is given out as a pattern, as a target. Then the schedulers try to live up to that target, dealing with real-life constraints as they come up. We just say, okay, you ought to take 50 percent of your cargo around the cape to Rotterdam and 50 percent through the Sumed pipeline, the pipeline through Egypt. That will be your most economical way to get it there. Then they do their best to stick by that schedule.

DOENGES: And this is like a general plan of operation for a whole year?

BISHOP: Twice a year for two years and once a year for five years. But where it breaks down is that there are three or four months before we do the review. If something changes, these guys are still scheduling their ships based on the old pattern. Also, they panic. For example, if a ship is tieing up a pier because it lost a pump. Now, all of a sudden, there are two more ships backed up and it ripples through the schedule. So people quickly make decisions based on experience but not necessarily with economic justification behind them. For example, "we will get this

one out of here, load it, and we will send this one to another port." After a while, it gets back to steady state.

But I am not so sure that in the interim they are really doing the most economical thing. I do have a feeling that they are getting very good at it. Our people have a lot of experience. But if you send a ship in the wrong direction for one day, that is \$30,000. Ships are worth, even in a depressed market, \$5,000 to \$10,000 a day, and I have seen VLCCs going for \$120,000 a day over port and fuel. So little mistakes can add big bucks.

So our problem, I think, seems to be in the very short term. With a small number of ships and a small number of trades, I feel reasonably confident that we can allocate them in the long term fairly well. We can set the modal pattern. That is not a problem. It is dealing with the short-term problems.

JAIKUMUR: You mentioned that the ships get into a pattern and keep doing the same thing over and over again. Tom Baker mentioned yesterday that if either supply constraints dominate or demand constraints dominate, you might change strategies. You might do different things. I am trying to get a feel for when these patterns would shift, for instance. I do not know, maybe it is not the same problem.

BISHOP: He was talking about the way of using a model to schedule ships. We do not have a model.

MACLEAN: Yours is strictly a heuristic solution to a current problem, as seen by your schedulers?

BISHOP: Right. We have got a guy who was scheduling convoys in World War II. He goes down there and says, "let us all do this," and it looks like it is right.

MACLEAN: Has there been any attempt to document the process that he goes through?

BISHOP: I could almost ask Dave Ronen better than me. We have them under contract to take a look at our problems on kind of a low-key mode

to see what they can do for us. The last set of stuff that I had sent them was what we had documented, as best we could, on why schedules had changed. For example, here is the schedule as of the given date and then 15 days later, and it is altogether different. We tried to document the reasons why--the ship broke down, the cargo was not there, there was nobody to tie up the ship, all kinds of reasons. So we are in the process of collecting data on how you load tankers, what options you have. We are working on the problem, but not on a priority basis.

JARVIS: Do you have measures of utilization or efficiency? If so, how would you compare your manual methods' capability with, say, published standards from the industry?

BISHOP: As far as I know, there are no published standards. We do not measure efficiency in a way that would affect scheduling. We measure efficiency by counting delay times on the vessels, which is very difficult to do, particularly when delays start causing other delays down the line. Operations is always saying it is an engineering problem and the engineers are always saying, well, we were doing it because the ship was down for operational constraints, anyway.

But we collect the data and in most cases it is done on the basis of a target, e.g., it ought to take you a day and a half to clear this port. Then we go through and we see how many ships cleared and what was the average time. So it does not really affect the schedule. It is almost a port-performance type criterion.

MACLEAN: How does this performance evaluation feed back into the rescheduling operation?

BISHOP: It does not. It feeds back into the port performance.

MACLEAN: I see. So it raises criticism, but it does not offer a solution.

BISHOP: Well, they break it up between uncontrollable and controllable delays. And when you get enough controllable delays at a given source, you go deal with the source and hope it goes away. A lot of it has to do with lack of tankage capability--a ship could not discharge because there was no place to discharge, or a ship could not load because there

was no place to discharge dirty ballast. If this goes on long enough they come down and say that we must need a bigger ballast tank.

RONEN: May I make a general comment about standards in the shipping industry? In my experience, I found out that ship schedulers are usually very optimistic. They schedule the ships according to terminal speed. They schedule the time in port according to the pick-up rate. As far as I have seen, they hardly ever learn. They do not have feedback.

One of the reasons is that ships cost a lot of money per day, tens of thousands of dollars. They do not want to build slack into the schedules. About five years ago, I was involved in a study team that analyzed industry data and tried to compare, for example, the time that the tanker spends in port to what the industry standard is (which does not exist, but is perceived by certain people). It turned out that the actual time, on the average, was about twice as much as the standard. So we should be very careful in speaking about standards.

BISHOP: I also might mention that, and I do not know if this is typical of the industry, our schedulers are at least one step removed from the ports. You cannot go down the hall and ask a scheduler if some ship has sailed. He has gotten it to the port. That is the end of his problem. He has it at anchor out there. Now it is the terminaler's problem, and he does not even talk to the terminaler. He talks to somebody else in another supply and distribution department who talks to the terminaler who talks to the pier. So they are not even aware of the problems. Their problem is to get it out at the anchorage and then they forget about it. The only time a word comes back is when they say do not send me any more ships because we have a full anchorage.

DOUGLAS: David Ronen's comments were fairly accurate, but I do not think that the negative aspect of them is valid because a scheduler has to be somewhat optimistic.

RONEN: I did not say it was good or bad.

DOUGLAS: I just wanted to take that out of it because of the significance of the ship coming early, which it will do sometimes, and then you are probably in a lot of trouble.

RONEN: Right.

DOUGLAS: Whereas, especially the way it has gone, if you say it is going to be here on the 25th, okay, it will be here on the 25th. It could be the 26th, too. Maybe even the 27th. Who knows? So, it is a comfortable feeling and a comfortable situation to have.

RONEN: I was just commenting about the practice that was going on.

SCHRAGE: So you are saying the system just expects things to be late, but not early. It would cause a lot of trouble if they gave the expected dates.

DOUGLAS: That is right. For example, if you agree to carry a cargo for somebody a month from now, the lay days might be from, say, the first to the 15th. But your ETA is perhaps the second. This gives you enough slack so that if it is delayed some place because of ice in the port or because the coal is frozen, etc., then you still have a fair shot at it; that sort of concept pervades the whole system.

MACLEAN: This would be fine insofar as chartered tonnage is concerned, but would you run your own tonnage in the same way?

DOUGLAS: Yes. I do not see the significance of the difference.

MACLEAN: In one case you are working with your own asset. In the other case you are working with somebody else's asset and you merely have a contract for service.

DOUGLAS: Right.

MACLEAN: The economics, I would think, would be somewhat different, depending upon whether it was your own ship and whether you were able to move it faster and get it on out.

DOUGLAS: There is really no magic to it. Just because it is yours does not mean to say that you can push a button and get it there any differently than the operator who is coming in on a contract. There might be a circumstance where if it is your terminal and a couple of

your ships are in line, you might bring one in ahead of the others because of that sort of situation and where you have too many lay days, say, down at Norfolk.

But basically, there is no magic. If there is going to be a storm that is going to delay that ship a day, it is going to be there.

MACLEAN: Yes.

DOUGLAS: So that same concept pervades if it is somebody else's ship coming in bringing our cargo. We might say, we want it between the 1st and the 15th. And if his ETA is the 13th, we might say that ship might not make it. Then we make our judgment from there--either accept the ship or say, yes, we will accept it, but is that a good ETA and how firm is it? Some people are a little fuzzier than others as far as the information is concerned.

MACLEAN: Would you have any ability to speed up the vessel? I mean, after all, the nominal speed of the vessel typically is a little less than the design speed, because you usually do not run it up into the extra power necessary to drive it through. You may be able to make an extra five percent on the speed if you pushed it.

DOUGLAS: That is possible. But unless you have a fairly lengthy run, it is not really significant. Moreover, with most merchant vessels, if they are going full speed, that is pretty well flat out. If they are slow-steaming, then you do have some flexibility there and there might be some playing around with a few hours here and there to speed them up to get in ahead of somebody else and get in line.

But getting back to the days and being optimistic, it is not like a streetcar. You come in with something less than that.

COPPINS: You do not gain anything by getting it there early.

DOUGLAS: Only to get in line.

COPPINS: Yes. And the idea that you are going to gain because it is your asset and you are going to avoid inventory on it probably is not that valid. There is enough slack that even when you have a load,

it is going to sit there for a while.

I did some work with Reynolds Metals over the summer and in many cases they were shipping their own material around; it was bauxite and aluminum. If they got it there early, they would just add it to the pile. The only times when they really had to be careful was when they were supplying the plants, and the plants had very limited inventory possibilities. But if they got there early, they might not be able to unload it. So they would end up paying demurrage.

They did have limited inventory capacity at the plants and so they really had to get to the point where they could unload the ship completely. Many of the ports had no facilities for intermediate storage of this stuff so the ship could discharge and leave.

So there is actually no benefit to getting there early. There is no push on them. What they really want to do is get there, if they are supplying the plants, within the window so that the plant does not run out of material. Obviously, that would cause a big problem. But as long as they keep the plant moving, it does not really matter, insofar as it does not mess up the schedule down the line. They patch the schedule every morning, anyway. So it is not like publishing the schedule Monday and it is good for the next two weeks. Rescheduling is the first thing that the guy does in the morning.

Not only is there no incentive to get there early, there is a disincentive in many cases. I think that is one reason that ships are scheduled the way they are.

BALLOU: That is for the POL side, too, for pick up?

BISHOP: Yes, if they give us a window, we can get there early. But if you get there two days early, you have lost two days. The odds are that you are not going to be able to discharge. If you get there two days late, you are not going to hurt anything.

RONEN: If you get there early, you are probably burning more fuel than you should be burning.

MACLEAN: How does this stack up with the problems that you have faced Tom Baker?

BAKER: Well, in the problems they describe I can see similarities to problems at EXXON. They are a subset of all the problems that we look at. There are certainly a lot of scheduling problems that look more like straightforward assignment problems. When you reduce your control of the inventories to just a volume and window, then you do not worry about the inventory manning problem. That is somebody else's problem. You have gotten your ship there with the cargo or pick up the cargo during that window. Well, depending on how many sources and sinks you have in the system, it could end up degenerating to a straight assignment problem. Given the right supply conditions, you might be able to say, you are going to take this carrier and put in there, or this tanker and put it in this service on a fixed basis right on schedule, etc. Those problems are fine if they degenerate to that point. There is not much of a scheduling problem to worry about. But anything could happen along the way.

But you can go back to the other end of the spectrum. There are problems where you have got a ship that may not return to the same service for a long, long time. And if you are really going to try to schedule that kind of operation, you have got a traveling ship and there is no way to get rid of it--and so the problem now that you are dealing with is really a problem. We have got problems of that kind.

So I guess I can see parallels in management problems. I cannot see any generalities in terms of what we can say about the industrial scheduling problems.

MACLEAN: It appears as though the process is sufficiently segmented that there is not much feedback from one segment to the other, except if there is some post-voyage analysis of a particular fleet unit or a particular port. But then there is no coordinated structure to this. Who sets windows, for instance? If a refinery says, we need to have something here between the 1st and the 3rd, is this strictly their function or is there some other responsibility being discharged by them that has been assigned somehow else?

BISHOP: It is their function as long as it does not cause a corporate-wide problem. If, for some reason, it does, then we are going to complain. But as long as they play within the plan that they have told the corporation they are going to play, as long as that scheduling, that window, falls generally within the plan, we just fill the request.

BALLOU: Is the issue perhaps a linkage between long-term and mid-term planning and what you were saying is the refinery requires short-term scheduling? Or is there long-term scheduling that is done as well?

BISHOP: We have tried it. It does not work. The ships get out of synchronization too fast. By the time you look at one or two voyages the schedule is off. I did not mean to imply that our ships run from the Persian Gulf to Rotterdam. We do have one ship that does that; it may go down to Rotterdam and then go to Sidi Kerir or over to Freeport. But for long-term planning, we just say certain ships will operate in certain trades. In the short term, they are all over the place.

BALLOU: Do you see any linkage between them?

BISHOP: Well, I was going to say that there is always this problem that if you go from a schedule on a lifting basis to a schedule on an average-runs basis where you take some sort of average and you say these ships will be out there and this is a scheduled basis--then for some places they run in together. And where they run together, it does not make sense.

Now I think we have a break that the military probably does not have. We can schedule out long enough in time to where we can put that break where we do not really care. We solve the short-term problem of getting the known requirements where they must go, and that keeps moving up so we keep this interface from causing us an operational problem. Unless Saudi Arabia goes down or there is a real major disaster, whatever is going to happen in the interim, in this phase that we cannot define, is going to be pretty much like what we are seeing today. There would not be a radical change.

If there is going to be a radical change, we do not know about it anyway, so we cannot plan for it. So we essentially make the assumption that we can keep approaching this thing and solve it as we get close to it. And the only reason that we have to worry about the long term is to set our requirements for the number of ships and the number of tanks and make sure that we can fulfill our own requirements.

So it is not really a scheduling problem. Our scheduling problem is only short term.

MACLEAN: So you have got a long-term assignment problem which really allocates a ship to a general service.

BISHOP: Yes, but we do not have a long-term scheduling problem. We either use essentially financial analysis with sensitivities or else simulation. We try to attack the problem that way since, in the long term, we do not know what our requirements will be. I am sure that EXXON will agree, that the industry as a whole, and we are no different, is very bad in predicting what is going to happen five years from now or even two years from now.

We have gone back and plotted what we thought was going to happen and first it goes one way and then it goes another way. There is no point. I mean, if you do not have any data to work with, there is no point in being sophisticated about it. There is no such thing as having precision with no accuracy, you know.

But in the military, from what I heard yesterday, I guess they define a situation. It may not be real, but at least it is something that they can plan for. They have so many tanks and divisions that they want to move at a given time, whether that will ever happen or not. Taken as a contingency, i.e., if you say this is going to happen sometime in the future, what is the optimum way of doing it? Then you have a problem that you can work with, you can drag out the solution if you find another problem that looks like it, I guess.

BALLOU: I think your comments on the precision versus accuracy are appropriate.

JARVIS: With regard to the driving force, I am curious to know something. First of all, let me make sure that I understand. I believe you said that you are almost totally manual at Chevron.

BISHOP: Yes, we have some computer assistance in long-term planning.

JARVIS: With regard to driving force, what would happen to cause you to reassess your scheduling mechanism? How would you know when you are due to make a change?

BISHOP: Normally, we reassess the situation every four months. That is our policy. In the interim between that, we hope somebody catches something. It is that simple. I will sit down in my office and I will see the tanker mark has come in at W-20 or W-22. I kind of remember that we set this thing up at W-30. I am sorry. I guess that is a term that you are not familiar with.

MACLEAN: No, it has not been mentioned yet.

BISHOP: It is a nominal system for chartering tankers. There is a World Scale 100 and World Scale 20 is 20 percent of that rate. There is a big book that has all of the World Scale 100s and it is set up so that for any percentage of World Scale, the operator gets the same net back in any trade. So you do not have to argue dollars-per-ton; you just argue percent of World Scale and the higher they can get, the better off they are.

So if we are planning on a higher payment to independents to deliver our cargoes, and we really see this come to pass, it might stick in my mind that we had better take a look at this again or get somebody to look at it. Or if we planned on taking 300,000 barrels a day through the Sumed pipeline and all of a sudden they say you can only have 250,000, then we have to reassess the situation. There is no set procedure. We hope that one of the people in the cycle will recognize that something has changed.

MACLEAN: Perhaps it is worthwhile to define what World Scale means at this point. World Scale 100 used to have a value, back during the post-war years. When T2 tankers were a primary transportation medium, there

used to be a nominal value. As the tankers got bigger and bigger and their unit costs for delivery kept going down lower and lower, this basic reference was never changed. Consequently, if you look in the charter market and find out that tonnage had been chartered at World Scale 22 or something like that for a run from the Persian Gulf to the east coast, for example, you immediately go back to the scale and you can find out how much that is costing and how that fits into your marketing as to whether you are paying more than you should or whether you have got an advantage.

SCHRAGE: Well, is it still done that way?

MACLEAN: Oh, yes.

SCHRAGE: Why?

BISHOP: It facilitates the bargaining for ships. Suppose there were no standard reference and I am an owner. I have got one ship and two people bidding for it. One wants to go to Japan for \$10 a ton and one wants to go to Rotterdam for \$18 a ton. Which way do I go? I would have to have an analyst there to figure out where I was going to make more money. With World Scale, all I have got to do is get the highest World Scale I can and make that choice. It is not a perfect system, but it works out pretty well.

DOUGLAS: In the World Scale 100, there are different rates for each trade.

BISHOP: Right. Any percentage of World Scale gives you the same return after port and fuel costs on a dollar's pay basis. And you can collect premiums. Everybody knows that if you want to go east, you are going to have to pay a 5 point premium. So W-25 east is the same as W-20 west.

DOUGLAS: They change on sizes, too. A 200,000 ton cargo might ask a different rate than a 100,000 ton or a 50,000 ton in so far as the level of World Scale.

MACLEAN: And then there is also the question as to whether you take a full cargo or a part cargo, because if you have a big ship and it is

being offered at a very favorable rate, it may be advantageous to you to use that ship for less than full cargo because even though you will be paying for it, it still may be more advantageous in the given time than getting the right size ship for it.

JAIKUMAR: How is World Scale determined:

BISHOP: An independent brokers panel in London puts it out twice a year and updates it intermediately if there is a major change.

MACLEAN: It reflects current costs?

BISHOP: Current port and fuel costs. It is a handy device when you consider 500 to 700 spot charters are made each month. You would have to have a lot of analysts to support that much negotiation. The dry bulk trade is different. It has its analysts. It is a much more complicated situation to charter a dry bulk ship than it is to charter a tanker.

DOUGLAS: Dry bulkers do not change the ports as much as oil tankers do. And this way, an added advantage of World Scale is that you do not have to pin down exactly where the cargo is going to go, or unload, really. It is just an area to an area. You do not have to negotiate a rate for each possibility. The World Scale has a rate for each of the combinations and you just negotiate the World Scale for the voyage of interest.

SCHRAGE: Can you negotiate a World Scale rate for shipping and then change it slightly to take advantage of it, because the rates are better for you?

BISHOP: The fixtures usually are not made port-to-port; they are, e.g., PG west, which means west of Suez, or UKC, United Kingdom Continent, or UKC option Caribbean. They are general broad areas and the same rate applies to Hamburg and to Rotterdam. You just have a different flat rate in the book by which you multiply the percentage, and then count the difference in lays and port charges, etc.

SCHRAGE: So do you specify when you do the negotiations that you would like it to go to Rotterdam, but you really are thinking of taking it to the Caribbean because the Rotterdam world rate is not so good for you, whereas the Caribbean rate is good? And the carrier says, "yes," based on thinking it is going to Rotterdam.

BISHOP: There is not that much difference. In most of the west trades that we deal with, there is perhaps half a point, which in today's market is not very much. If you are playing that kind of game, the carrier will charge you for the option of going to the Caribbean, e.g., W-20 for UKC and W-22-1/2 for UKC plus Caribbean option. People get pretty sharp at it; it is their livelihood. It is pretty hard to "game" on somebody. There are not too many amateurs in the business.

MACLEAN: I wonder if there are other things that we should talk about with respect to what the process actually consists of? What kind of a process do you use in your scheduling, Burnie Douglas? Is there some systematic structure that you have or is it, again, a manual process? I understand that Bethlehem Steel has gone into some computerized scheduling. Would you elucidate as to what your process is?

DOUGLAS: I can characterize it as a real fast calculator. We have the ability to run these schedules out by the scheduler very quickly. And if the scheduler can schedule his fleet of vessels out very easily and conveniently and, as I say, with a real fast calculator concept, come up with the results, that is 80 or 90 percent of the problem.

For the longer term look at assignment and that sort of thing, we have a linear programming model that looks at the next year or the next five years, or however many years one wants to go. There is a link coming that will tie that linear programming model into the fast calculator scheduling system. But the user is still very much there and in control.

MACLEAN: Is it interactive?

DOUGLAS: Yes.

MACLEAN: Would you classify it as a computerized heuristic solution?

DOUGLAS: Could I ask again for a definition? Does heuristic imply a trial-and-error approach.

MACLEAN: You try until you get a fit. Now you have no idea as to whether it is a good fit or a bad fit. It is a fit. And then it becomes an upper bound against which you can work if you want to try to get a better fit.

DOUGLAS: Yes. We have been working with a manual interface between this assignment concept and the manual scheduling process. And in that context, it is a heuristic in that the scheduler uses his experience, with various inputs from corporate headquarters as far as what sort of percentages of the cargo should be carried back-haul, what should be carried just back and forth in ballast, etc.

Then this tie-in that I referred to is coming, and it is strictly a rounding-off type of thing. If the LP says 4.6 voyages, this tie-in is going to use a lower bound of 4 and an upper bound of 5 to start off with and see how that works.

So it is not optimized in the usual sense of the word. But it does use an output of an optimization routine to do the scheduling, with the user's help. Now let me emphasize that. The scheduler is right there and he has control over how far out he wants to fix the voyages of a vessel. By "fix" I mean not contracted out, but just in his mind or in the operating people's mind.

NEVEL: What is the long-term use of this program? You said it is for long-term planning. Is this for fleet-sizing?

DOUGLAS: Basically, it is for fleet-sizing. We can take a look at the consequences if we contract for a million tons of coal. What would that do to us? How many vessels do we need if we assume that 75 percent of the trips are going to be back-hauls, or 50 percent, or whatever? It is a fleet-sizing, cargo-sizing tool.

MACLEAN: So it is an operational analysis tool, then. You exercise it to see if you like the way the full scheme works, but you do not use it for the actual fixing of the ships on their specific voyages. This is essentially an input as a suggestion of where you might start.

DOUGLAS: Yes.

SCHRAGE: What are the decision variables in this model from this LP? Is it ships taking a particular route in a particular week?

DOUGLAS: No, no schedules. It is on the time period; typically a year, where if we are going to have x-million tons of ore and we assume, say, so many million tons of coal or grain, etc. And if we assume a certain fleet size, how does that look? Can we cover it? Are we going to have excess capacity? Should we be looking into chartering in another vessel? It is that sort of thing.

NEVEL: How often do you do this? Is this an annual review?

DOUGLAS: There is a little ebb and flow as far as emphasis is concerned. It is at least once a year, and sometimes twice a year as far as the five-year look ahead.

Let me just interject something here because I think it is rather significant, particularly in view of the discussion yesterday. Sometimes it is difficult to assign values to the importance of, for example, getting a certain division there at C plus 5. Obviously, everybody recognizes that it is difficult, and yet, there are some analytical pressures to say "give us a number so that we can come back to you with the answer."

I think it is important not to look at that as the analyst coming back to the operating people saying "I got you an answer." It is much more important to put that back on the operating people in saying the following: "If you want it there at C plus 10, this is what it looks like. And it is going to cost X bucks. However, we have run it so that you slacken things a bit and it can be there at C plus 15, with perhaps a 90 percent probability that it will be there at that time,

depending on where the vessels are when you start. But it is going to cost you a lot less. You, Mr. Operator, Mr. Chief of the company, how do you want to do it?"

MACLEAN: You are giving him information that can be put into a value judgment.

DOUGLAS: That is right. And many of those things have to be value judgments. In the commercial area that was described yesterday, if you want a sailing every week from this port, then this is what your profit will look like at the end of the year. However, if you want to slacken it to once every two weeks, with whatever amounts of cargo that generates, then your profit is going to be in this range.

You, Mr. Chief Operating Officer, how do you feel about your customers? Will they give you as much cargo if you are there every two weeks? But the profit is still a very handy number for him to digest and see what he thinks.

MACLEAN: But again, this goes right on up to the top. The strategy, the thrust, has to come from the top as to how they view their information requirements and their decision requirements for their corporate management. And this, apparently, never gets down below them. Is that correct? Now, you do not feel anything but the segmented outcome of this at the working level, essentially.

DOUGLAS: I do not think I said that.

MACLEAN: I am trying to see whether I am reading you correctly. Because it seems to me that if you are feeding quantitative information up for them to use in making a decision and you do not know what that decision is going to be or on what it is going to be based, that is a strategy which they have not exposed beyond that level.

DOUGLAS: Well, there has been some discussion such as, "hey, we want a sailing every week."

MACLEAN: Okay, this becomes an operating strategy.

DOUGLAS: That is right. And to the extent that that could be flexible, it might come down, "hey, it is wanted every week, but what does it look like every two weeks" - that sort of thing. Obviously, there has to be some interplay between the working level and the decision level. But my last comments were not specifically with regard to my situation; they were more general, as to the interface between analysts and operators.

JAIKUMAR: I have a question. What you are essentially saying is you are looking at a set of scenarios. For each scenario, you are defining some optimal strategy. If the model is sufficiently complicated, you cannot really work with a large set of scenarios. It is a finite and a reasonably small set.

How does it get defined? Do you view that, the defining of the scenario itself, as an analytical function. Or do you think that that comes from somebody else--from stock management, perhaps--who defines this as the scenario and says "now go ahead and do it."

DOUGLAS: Well, the tools are there for the analysts and the middle managers to play with so that they can get a good feel for the different scenarios. Historically, it used to be when it came time to look at the annual "what-are-we-going-to-do," all of these scenarios would have to be shifted down into one. And with a hand calculator, it took over a weekend or so to put all the numbers down and come up with a number.

Well, you just do not have that same problem if you can play with a scenario and come up with an answer and then try it again with a different scenario, etc. You know, you do not have to take them all. But as long as you can play with them, that is the key.

MACLEAN: I would like to ask a question now that should try to bring some of these ideas into focus. What should be done in the way of developing our capability to better handle these problems that we have been talking about? I mean, your problem is largely manual because you have got some very capable and experienced people there. If you lost those people and did not have any means of readily getting somebody

of equivalent capability, what would you have to do to keep your operation going?

Would you try to automate some of it, get somebody in to see if you could not routinize it, computerize it? What do you think that you would be doing to try to extend the work that you have done into a more computer-faster solution capability?

Tom, where do you see EXXON going in terms of trying to get more rigorous and tighter control using a simulation capability or algorithm development to treat the scheduling problems that it has to face?

DOUGLAS: I think we are pretty well squared away in the scheduling and assignment areas. The area that we will be concentrating on is what you alluded to before in terms of feedback: how well we did against what we estimated. That has been somewhat lacking for a long time.

BISHOP: I am not worried about losing schedules.

MACLEAN: You are not?

BISHOP: We have got a kinds of schedules. Some are better than others. And they do not all go at one time, as planned. I guess what we are a little worried about is that maybe we are leaving some dollars on the table by not having a computer-assisted human.

We are looking at alternatives and we may find something that will work. It just has to be cost-effective.

BAKER: I think in general the first thing you need is the on-line information system. The fact is that most of our scheduling functions, even though they are done on worksheets with stubby pencils, are assisted by some kind of on-line information system.

We have got one or two, I guess, examples of scheduling systems that have been more automated. Unfortunately, not all the other existing on-line information systems are in a form that will allow you to do any kind of informational decision-making, or even algorithmic decision-making.

I guess what we are trying to do in a number of areas is move towards information systems that will support that type of function. We do have a lot of situations where we have one scheduler, e.g., in a profit supply environment such as coastal supply. He is really the only person who has come up with a good tip in terms of information.

MACLEAN: So this really offers the opportunity to gather the information and systematically have it available as input and for interactive judgments with respect to refinement of the specific schedule, or as an input for another automated approach to an optimization of the schedule?

BAKER: That is right, short term and long term. If you have got it set up right, you can use it for all of those, and reconciliation too.

AD-A123 464

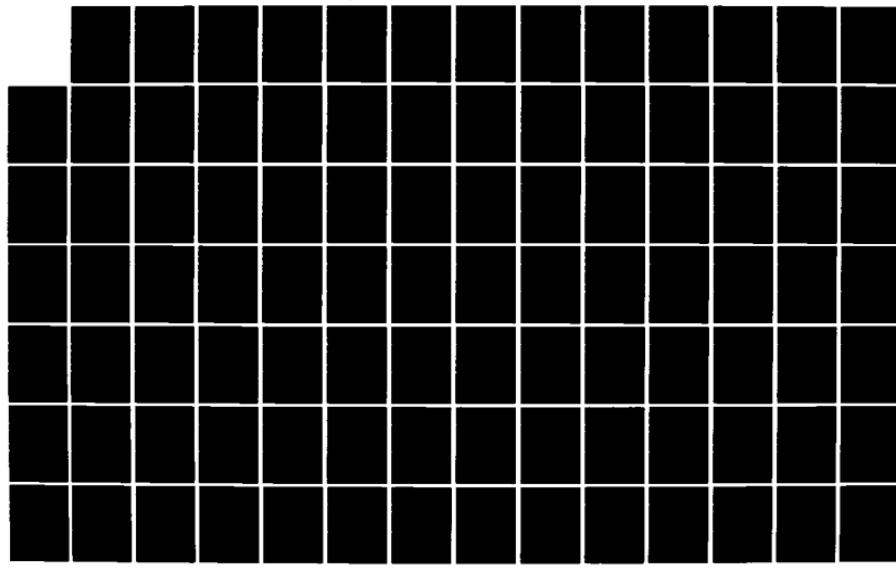
PROCEEDINGS OF A SYMPOSIUM ON CARGO SHIP ROUTING AND
SCHEDULING HELD WASH. (U) GEORGE WASHINGTON UNIV
WASHINGTON DC INST FOR MANAGEMENT SCIE. R M SOLAND

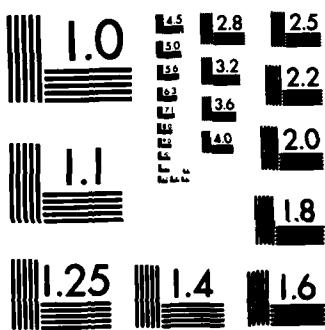
3/4

UNCLASSIFIED

15 DEC 82 SERIAL-TM-69258 N08014-82-G-0016 F/G 15/5

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

Discussion Group 2 - Liner Operations

Leader: Donald K. Pollock
Maritime Administration
U.S. Department of Transportation
Washington, DC

POLLOCK: Members of Discussion Group 2 on Liner Operations, you are asked to consider the following questions:

What are some of the problems that carriers face or requirements that liner operators foresee which could be dealt with through the use of models or optimization techniques? What use is now being made of available methodologies? What new optimization, modeling or analytical techniques are on the drawing board? When will they be available and what types of problems will they be useful in solving? Are there specific projects that industry or government need to undertake to improve liner operations?

This session will not include discussion of two important areas of liner operations. These involve the weather routing of ships and terrorism. However, Jim Mays, an expert on ship weather routing, who is in attendance at this session, has indicated that knowledge of the weather routing of ships can be a useful tool in developing ship scheduling algorithms as well as in the maintenance of existing schedules.

Let us now go to the liner operators and put them to work right away. Would you please tell us something about your company's operational problems and what models, optimization techniques or scheduling algorithms are used to solve these problems? If not, do you foresee a need for the use of these techniques in your operation.

CONWAY: No. I guess I will repeat a little bit of what I said last night. To us, a pure container operator, the schedule is a simple thing. We have a weekly sailing. It is in multiples of seven days. We do not vary on port calls. We do not vary on port rotations. So, the schedule is pretty well set.

Without any formal models we have done calculations that we can use to make sound business judgments. One of the things you have to do is sail a ship every week, and if you happen to be late because you are late getting through the Panama Canal or something like that, you have already done the calculations to know that it costs about \$40,000 to pick up a day, 24 hours, on a leader class container ship going to Europe. And that kind of thing helps you to make some basic decisions.

At least as far as U.S. Lines is concerned, it is not much more sophisticated than that. We do use models for stowage, but not in the scheduling of our ships.

BRIERRE: We do not use anything particularly sophisticated, either. I think we have a larger fleet than U.S. Lines. Thirty-eight of our ships are break-bulk ships, so that the schedules are not necessarily set. But an individual, in conjunction with some other people, determines which ships go on which routes, by matching known or existing cargoes with particular schedules. He is our most complicated piece of machinery.

I think Russ Stryker, although he sounded pessimistic, was fairly accurate. When you consider the number of ships that you have to deal with, you cannot necessarily apply these algorithms and get the payoff that you need to compensate you for the time and effort.

With regard to our container situation, I am sure that our containers are on computers. We keep track of our inventory. We reduce the time that they are out of our control. We probably do the same thing with our barges. But ships -- no, we don't.

POLLOCK: Bob Temmler, from Moore-McCormack. What is your opinion?

TEMMLER: We do not use any types of models, either. We are currently sailing ten vessels. The ports of call are rather repetitive, and the only scheduling decisions might be to hold up, to skip, or to switch calls by vessels. The costing or the profit analysis is done on a voyage basis, rather than on any type of an annual profit basis. We use no scheduling models.

PENTIMONTI: I can tell you just briefly what our situation is at American President Lines. We operate 21 ships and have, in the past, put some models together. They are not optimization models, but more simply computational models that help give us data, once we have made decisions as to routing and scheduling, as to variations in the routes. They are operational tools. Such a model is also used as a planning tool. We simply put routes and schedules in, and use it to tell us, as a management tool, what the various routes that we select manually do to us, for the autumn line review.

That is about where we stand. We have some optimization models that we use for stowage, as you have mentioned. Those are surely not hooked into anything that looks like routing.

One thing I mentioned to Paul Mentz this morning is that we have an interest in the development of a bit more sophisticated model than we have now use for planning purposes. That will help us make some decisions as to feeder ships. Although we operate feeder ships on certain of our routes, and as you know, are going to be introducing some large new container ships soon, and we are looking at the utility of the optimization that can be obtained by using various combinations of feeder ships and larger line haul ships.

That is a "motherhood" issue. But because of so many of the other factors that come into play when you consider the feeder ship operation, we thought that we would be better off having a model of some sort to help us understand the various rankings of the various feeder ship scenarios. We do not, however, intend to make a model that is sophisticated enough to optimize a feeder ship-line haul relationship; we simply want one which, again, provides data once the manual selections of routes and feeder ports are determined.

POLLOCK: I would like to introduce you all to Paul Mentz, who has carried out a very exciting program in the shipping operations information system area. Paul, in your program, did you consider any payoffs or benefits that could be derived from these optimization techniques?

MENTZ: Over ten years of working with the carriers has convinced me that the words "optimization techniques" are really the wrong words.

The business is complex. It is changing, even though rotations tend to stay the same, and frequencies are certainly driven toward weekly and regular sailings. The business is changing all the time. Competitors are changing in the way that they add services or subtract services. Even the companies with relatively stable port rotations will look and say, "Should we call at Panama? Should we not call at Panama? What should we do with our Guam service?" These are things that have come up in the past.

But optimization does not seem to be something easily handled, discussed or considered. I think Gene Pentimonti put his finger on it when he said that what we want are computational models that allow us to essentially get rid of the drudgery of doing all of the calculations by hand, assess different decisions that we have in mind, different decision options, such as whether or not we replace a direct call with a feeder call, or a call in a new port location, or change the routing of our vessels. It is more that we will set out the framework of what we are thinking of, and then we would like to have computational tools to tell us what the result would be of implementing that framework.

And I am not sure I will go so far as to say "simulation models." I think even that is on the fringe of what is acceptable within the companies. But computational models, to be able to put in ports of call and cargo flows, and ship assignments, using actual experience in terms of data on costs and such, and then to have a computational tool which assesses the potential result in terms of net benefit or profit and loss for that sequence of voyages or timeframes—that is what I mean.

I guess what I am saying is that there is a need to go a little bit further, to have some improvements in these tools, but certainly not to the degree that was discussed yesterday by some, in terms of optimization techniques or even large models. I do not think that large models are acceptable to the operating companies.

As a separate commentary, we have worked with many of the companies for more than ten years. We have done work in the planning area. We have developed computational tools of different types. We have done work in-house at MARAD. We have met with a few modest successes here and there. We have had disappointments along the same route. And I think that the result has been a kind of evolutionary movement forward, in terms of capability and understanding and willingness.

I will conclude by saying one thing. There was a comment yesterday by someone that he did not think that the operating companies ever shared or talked about what was going on, or what might be going on, or what could go on in this field. And while that is partially true, I would like to say that there have been over two dozen meetings, on a national basis, of companies who are interested in fleet management techniques, with over 100 people attending each of them. And there was substantial discussion and exchange of ideas and information. So, to say that there has been no dialogue is not true. To say that there has not been enough is reasonably true.

POLLOCK: We have Bill Webster here who is a consultant with APL (American President Lines). Bill, would you discuss some of the things that you have been doing with APL.

WEBSTER: What I have been doing for APL actually is at the periphery of the things that were discussed here yesterday. And this is involved with the very specific, and I must say, also very difficult, problem of ship container loading. By that, I mean which container should go where on board the ship. Now, at first it seems like a simple problem,

but then the problem gets much more difficult as you dig into it. I must admit that APL has been remarkably patient in the four-year gestation of this whole system, which now is in use both in this country and in the Far East.

The idea is the following: A ship is a kind of funny warehouse, particularly if you are talking about containers. There are cells; there are vertical columns where you store containers; and there are slots in those cells. From a very simple-minded point of view, you have to worry about the following situation:

If a container for a further destination is loaded above a container for a nearer destination, the situation is called an "overstow," and results in an inefficiency in the stowage of the ship. That means that at some point downstream, if you load the ship that way, you have got to pull one container off and put it back on again. There are two cargo moves involved that will cost at least \$100, plus time. So, if you have a route, such as APL has for some of its Far East routes, in which you call at a lot of ports, avoiding the overstowage problem becomes extremely difficult. In fact, it is one that is not always possible to avoid.

But the problem turns out to be much more complicated than just overstowage. It turns out that the capacity of a container ship depends upon how you load it. In all of the discussions that took place yesterday, the impression was given that a ship has a fixed capacity. That is not true for a container ship. A container ship is often limited by its stability; that is, its ability to sail safely in the ocean. And that stability is affected in a primary way by the center of gravity of the cargo. So if you can achieve a cargo stowage with a lower center of gravity, you often have the ability to carry more containers. If you carry less fuel, you carry fewer containers--for the same reason.

Thus, there are all sorts of games you can play. The container ships are not necessarily displacement-limited. That is, you can often carry more fuel and still not go over the bounds of how much you can

carry. So there are games that can be played, and some of these games are certainly played outside of the system, called a CAP system, which I developed. Some of them are played internally to that system, as well.

There are all sorts of real life difficulties in dealing with container ship loading problems. Some involve certain ports, where you can use certain cranes, while in other ports you cannot use the cranes. There are dangerous and hazardous cargoes, D&H cargoes, that can go only certain places--if you want to go into different ports in the world. Refrigerated containers have to go either in a very well-ventilated space or on deck. Containers come in various sizes. There are 20-foot-long containers, which are an anathema to most shipping lines, and 40-foot-long containers. APL is going to be introducing 45-foot-long containers. Your ability to mix and match is very, very limited. On some ships you can store 20-foot containers in certain slots, 40-foot containers in certain slots, and either 20s or 40s in still other slots.

As a problem, it is extremely difficult. You have to look ahead. You have to know your trade several ports ahead. You have to have the geometry, which is not a trivial matter, specified for the ship, because it is a funny warehouse, and you have to have all of these other considerations ground in also.

The basis of our model--since we have talked about algorithms I should say something about that--is, in a sense, pure heuristics. It is a simulation approach. Several different loadings are attempted, using heuristics rules. The selection is based on a ranking matrix; that is, a multiple criteria, not necessarily linear multiple criteria, method.

The program is interactive because I firmly believe that in any shipping line it is impossible to crystallize all of the requirements and constraints into a computer program. And you need the person there too--you need to unburden that person from the bookkeeping of dealing with this difficult problem, but you still need to tap his insight and his experience. And the interactive mode is one that I feel very strongly is necessary.

So our CAP system is a simulation system. It is interactive and deals with this problem, which in a way affects the routing because, if you find your particular route does not lead to good stows all the time, you can change it.

POLLOCK: Thank you very much. It is time to give you model builders and experts a chance to respond.

PSARAFITIS: I was not here yesterday, but I have a question about the previous discussion. I want to ask the liner operators to what extent is the logistics of containers important to your companies. I have heard some people talking about this, but I am not really sure of the nature of that problem and what the lines are doing to solve it. Once you dispatch containers to destinations, how do you track them and dispatch them back to the origins and other customers; what about the logistics of all that?

PENTIMONTI: I can attempt to answer, at least partially. Although there have been a number of steps taken in our company over the last ten years to use and implement various computer-aided management tools for the logistics management of containers, we do have and have now settled upon a computer-aided system which helps us manage containers. Our worldwide service feeds information into computers on a common base from both foreign and domestic sources so that key management personnel in every location where containers operate are using a common based, computer-aided system which helps in the logistics tracking of the containers and management of the container utilization, turnaround times, the actual sources, etc.

It is a fairly sophisticated system, and I think much has been done in the past in other companies also. I think each company can speak for itself, but each company uses a computer-based model tool for its container management. Paul Mentz has been basically the grandfather of some of the systems which have been sponsored and developed by the Maritime Administration.

MENTZ: I will carry on. I hope the fellows from USL and Lykes will bear with me. If I say anything incorrect, please jump in and correct it.

We have participated with both United States Lines and Lykes Brothers at different times over the last ten years to develop this transaction-process-based tracking system. It is not necessarily a planning system, although there are planning influences and outcomes. Mainly, it is an event tracking system. It depends on people and rotations entering transactions into the central data base, which might be maintained anywhere, but certainly the data entries are presently made throughout the United States and presumably will grow to the areas throughout the world where these companies operate.

I would say that these event tracking systems are pretty well developed and exist in a number of different fashions in different companies. There are two systems that were developed in partnerships, as I mentioned, and there are reports available on both of those systems that are in the public domain.

POLLOCK: Paul, you raise an interesting question. A lot of what we have seen done, has been done with the government and industry working together.

MENTZ: Some of it has.

POLLOCK: Will this continue in the future? Is there going to be a problem for liner operators to obtain support for implementing new ideas? How has the MARAD R&D budget been faring these days?

MENTZ: I do not know if that is a main issue for this group, but I will just say it is not faring very well and leave it at that.

I would like to make a comment, however. Government-industry partnerships and planning tools are tenuous at best. It is not a comfortable way to work, mainly from the industry side. Suppose you get to make a real contribution to a company's performance and do it in a shared program with government monies; that requires some disclosure--not disclosure of proprietary data, but disclosure of the techniques

that were used. That certainly becomes uncomfortable to the sponsoring company. It does not want to share a truly successful project in the planning area. In other words, the better it is, the higher the level it is at, the more successful it is, and the more impact it has on the company's performance in a positive way, the more nervousness and discomfort there is.

So I would say that our work in the planning area has been successful to some degree, but it certainly has not met with a wide-spread common sharing and common development effort, nor probably can it ever result in that. I would say that there will probably be less work done in these areas in the 1980s than was done in the 1970s, even though the requirements within the companies in the 1980s probably will be greater than what they were in the 1970s. In the 1970s the government was trying to encourage and to promote, to pull the companies into these areas, and I think to a large degree the companies did move forward in these areas. Now, the question is: will the companies continue to move aggressively on their own? That remains to be seen.

WEBSTER: I do want to add one thing. When I first got into this business of container ship loading problems, I thought that it was probably going to be reasonably straightforward, using general algorithms tailored to a specific shipping line. It turns out that almost all of the effort is the tailoring, and what I, for instance, have developed for APL is a very nonportable system, because in order to make it work in the real environment there was so much of APL put in there that to apply it to another shipping line I am sure there would have to be some very major changes.

So, to add to what Paul Mentz has said, even if you want to encourage interaction among the shipping lines in order to develop something real and worthwhile, it may not be easy to cross certain boundaries.

KASKIN: I want to indicate one area of effort, not mentioned previously, to supplement the planning models that Paul mentioned previously.

There is a Navy panel, of which I am a member, and Paul Mentz originated it. We looked at a model that is a simulation model that Farrell Lines uses. It may be essentially similar to the one APL uses, but I am not so sure if yours is a simulation model or just a computational tool.

PENTIMONTI: It is a simulation model.

KASKIN: The model takes various routes, various assumptions, etc. It was an attempt to get more sophisticated, and Farrell Lines wanted to expand the capability of the model; they had a proposal in for it.

We at MSC have a simulation tool, also, but it is rather out-of-date and is not very much used in decision making. But it was probably state of the art when it was done about two years ago.

YOUNG: I want to address a question along that line to Gene Pentimonti. If you use a simulation tool, do you believe its results? Do you use the results of the simulation tool in your planning?

PENTIMONTI: It is relative. The results are relative to manually inputted runs. You have a certain degree of confidence that on a relative basis your data are accurate.

YOUNG: So if you were to have some techniques that would, in an automated way, generate alternative routes and select the best ones from those generated, would that be something that would be useful to you in using that simulation model?

PENTIMONTI: I do not believe so, and I say this because of the limited number of ships and ports and key variables that are available to us. Because of practical limitations the user can fairly quickly come to a listing of alternative sets of variables of rather manageable size. And I just do not think, with the levels of ships that we have in the ports that we serve, that any more automated optimization techniques for arriving at routes would be useful. That is my opinion.

YOUNG: It is a comparatively small problem then?

PENTIMONTI: That is right.

YOUNG: If you had an improvement in route selection that someone had overlooked in doing it by hand, would it have a big impact in terms of dollars if you then changed your routes slightly toward the optimum?

PENTIMONTI: Normally it would not. Normally, the differences are subtle when the changes are subtle. Obviously, the big changes in results come when you really shift the utilization of ships from one basic route to another basic route.

YOUNG: Sometimes the dollars can be very large, so that a one percent improvement can be significant.

PENTIMONTI: That is not the case here.

POLLOCK: Do you operators see any opportunities now, such as in the Slade Gorton Bill whereby you are going to be able to do cargo pooling and revenue sharing, in terms of giving you more flexibility to do scheduling and routing and make improvements in that area? Does anyone want to deal with that?

PENTIMONTI: I will take a shot at it from one perspective; I think this surely is going to be the thing of the future, in terms of the agreements that can be made to set up a small consortium or have pooling agreements. But the opportunities are going to be driven by the ability to develop practical relationships among the operators. I think that is going to drive the available resources.

POLLOCK: Rather than any analytical technique(s)?

PENTIMONTI: Rather than any analytical techniques. In my opinion, we are going to see the generation of those capabilities from the rather tenuous agreements among operators rather than from the analytic approach.

POLLOCK: What action(s) should industry/MARAD take next? Is there anything you can see that we specifically need to do that we are not doing now?

YOUNG: It seems from the users that there is very little that they are looking for. Let me pose another question. Bill Webster has developed a model for APL that does some kind of optimization of cargo loading. Do you find that useful? Does it save you money?

PENTIMONTI: Absolutely. It has been in place for about a year and a half.

YOUNG: Is that sort of a one percent type of improvement over what you were doing?

PENTIMONTI: Even if it is only one percent improvement, one percent on a larger ship is about 14 containers. And what Bill says is very true. The container ship really does not have a fixed number of slots. It really does depend on the center of gravity of the cargo mass, and the lower you can get that, the higher your capacity is. In fact, being slot-limited and deadweight-ton-limited, when you add ballast the drops in cargo CG (vertical center of gravity) are very important to you.

With respect to the extent that we can demonstrate that this system has given us that one percent or that fractional improvement, it is very difficult, but we surely know that we are making a much better attempt at optimizing now. The payback is really there.

YOUNG: So, for well-defined problems, you are saying that small improvements, using mathematical techniques, could be justified?

PENTIMONTI: Absolutely. It is just that in the area of routing, other than the weather routing, which was addressed earlier, I think accurately, we just cannot see the techniques that we use now giving us any improvement.

YOUNG: What features make the container stowage problem more useful, or more amenable to mathematical programming?

PENTIMONTI: The fact that you are talking about 700 to 800 containers; is that the maximum size?

WEBSTER: Yes. It could even be 1000 containers.

PENTIMONTI: Up to 1000 containers, with seven or eight discharge ports as a minimum.

WEBSTER: A minimum for the longer routes. You can go up to 12 or 14 ports.

PENTIMONTI: You look at those variables and possibilities plus all of the restrictions that Bill Webster mentioned, with D&H cargo, different sized containers, the port rotations. It is a very complex problem. And our planners, although they did an adequate job, often because of last minute cargo changes and so forth, were not able to make changes quickly enough in order to sail the vessel on schedule. So, I think the tool actually has two benefits: it does things in a more rigorous way, considering all of the variables and all of the restrictions, but it also does it quicker when there are changes and last minute adjustments in cargo.

WEBSTER: I would like to add to that. There is another, rather subtle, benefit to using a system like this also. To explain it, let me first back up a bit.

At any given port you can always make yourself look good at the expense of downstream ports looking bad, and if you have individual planners who are involved in this operation you can have a problem. So one subtlety to this is that it regularizes your approach throughout the system and you do not have instabilities caused by a personnel problem.

YOUNG: It would seem that all of these considerations would add to the routing and scheduling problems.

CONWAY: I would add to the discussion by saying stowage is a constantly changing problem, and the schedule is not; and that is why some kind of model is of great help to deal with a constantly changing stowage in the vessel as it goes from one port to the next along the same constant route.

YOUNG: So a feature of this particular system is that it reacts instantly to the problem as it comes up. It is essentially an online system.

CONWAY: To perhaps put it simply; stowage is an ongoing problem and scheduling is not.

SOLAND: And that is similar to the problem that Jai Jaikumar talked about yesterday. It is a problem that comes up repeatedly every week or every day, but with enough minor variations that you have to have a new solution each time, and the problem is big enough so that the combinatorial possibilities preclude doing anything by hand.

KASKIN: I think, given that most of the liner companies are limited to the subsidy restrictions, there is not much choice. There is not much for a decision support system to tell you other than minor things that can be done manually just as well as automatically.

We at MSC have about 23 dry cargo ships that we can use in liner or nonliner service. We have flexibility. We do not have any operating restrictions. It might be useful for us to see whether, for various service levels, a liner type of trade could be improved. We can change on a day-to-day basis or a month-to-month basis.

YOUNG: So you make changes, plus you have the complexity.

KASKIN: I can see that in the planning areas a company may want to evaluate changing its business approach, maybe even in the capital budget area of assets versus various trades. If they have good enough marketing data they may be able to use a simulation model.

PENTIMONTI: We use them, but even in a changing environment, or I should say in looking at all reasonable possibilities, we still can do it manually, and get the results from the computational model.

POLLOCK: We have got just a couple of minutes left. Are there any other summary comments that anyone would like to make?

WEBSTER: I would like to make one comment amplifying on what Gene Pentimonti was saying. If you get 10 more boxes on a ship from the Far East to here, that is maybe \$30,000. If you have a ship leaving every week it does not take very much arithmetic to see how much money you can save. That is an optimization problem that has a big enough payoff to be interesting.

VOICE: As far as container operations go, which you seem to be concentrating on, we in the military have the problem of moving all the cargo that is offered to us. We do not have any choice.

POLLOCK: Thank you all for participating in the session.

Discussion Group 3 - Crisis
Environment/Time Sensitive Scheduling

Leader: Joseph F. Ballou

Military Sealift Command
Department of the Navy
Washington, DC

BALLOU: I want to start by further defining the problem and providing a little structure, at least the way that I view the problem from the operational aspects. It may or may not be the same if you are talking about an analysis of scheduling aspects. What we hope to pick up here are some of the techniques that have been used previously. What are some of the approaches that we should look at? That is, what we hope to get out of this session is where do we go now?

As Captain Scott mentioned yesterday, it is a three-phased problem. The first phase is shown in Figure 9.1. I use the word "notional." Commercial people say you do not know what you will be seeing a year from now. With us, the time frame is, of course, much less than that, but we still do not know what we are going to be seeing in the way of real lift requirements. The information that we get, for instance, is that some Army or Air Force unit is going to move. They do not bother to tell you how many tons or whether it is rolling stock. You hope you can get some rough idea of the tonnage.

I have indicated some of our assets. Below each input, the kind and the number of ships give you the flavor of the problem.

We are legally bound in the way that we use our assets. Obviously, if we have ships under our control, those are the ones that we consider first. We are talking about approximately 23 dry (23D), and 27 tankers (27T) that we have under our control.

FORCE SIZING DECISIONS I

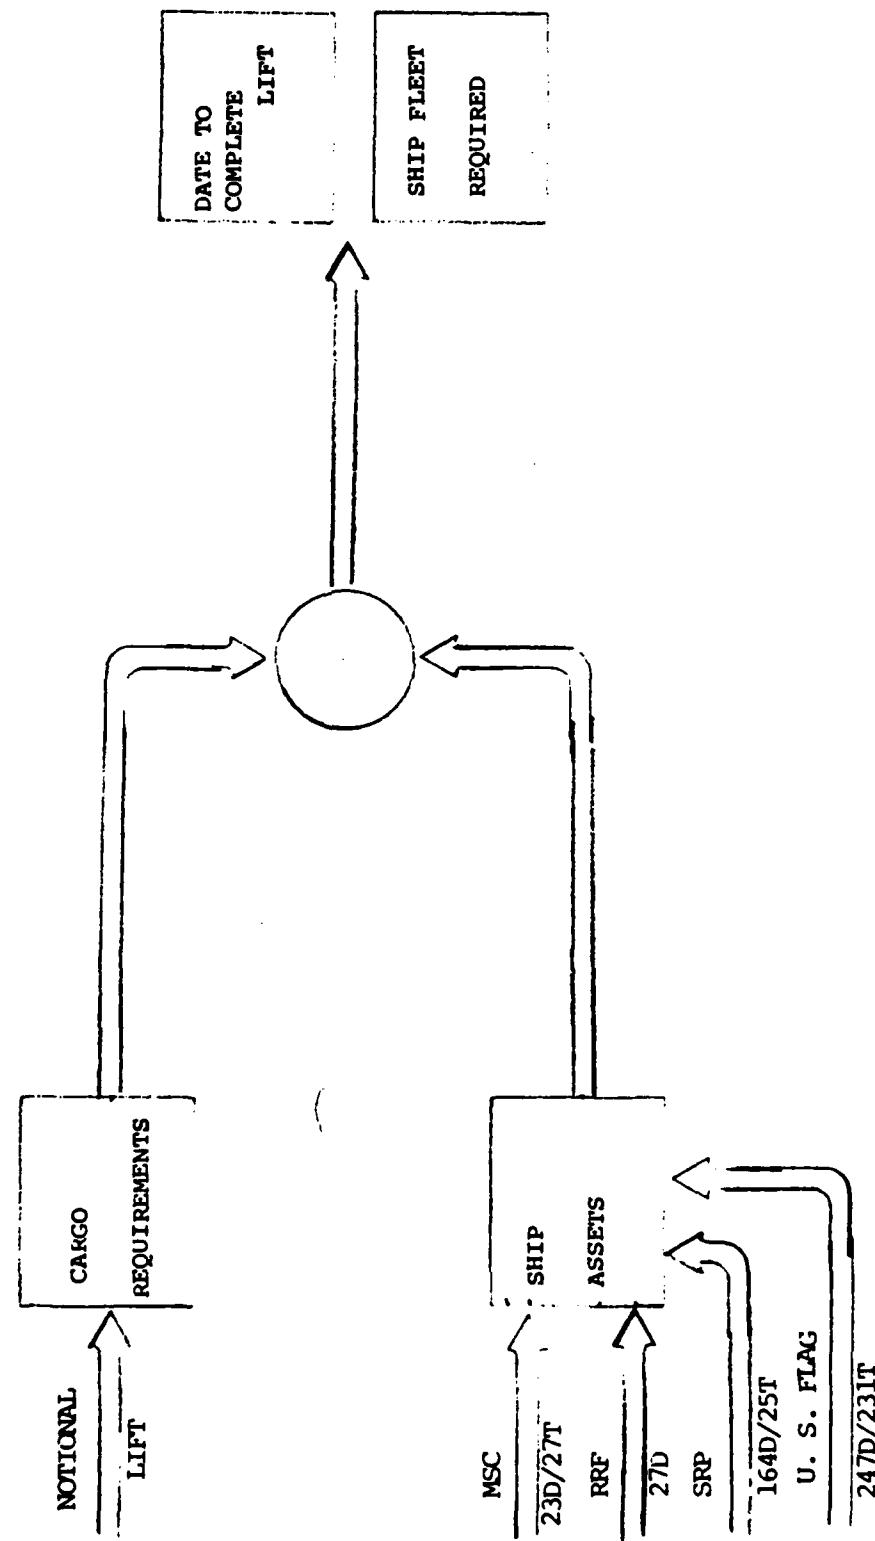


FIGURE 9.1

They may not be positioned to do the desired lifts. If you decide that you need more ships, you go to the market place and look for voluntary charters, either the dry or suitable tank-type ships. That is not explicitly put in here, but that is really the first step.

Some of you that we deal with get phone calls every once in a while. We do that to get an idea as to what might be available.

The Ready Reserve Force (RRF) has about 27 dry cargo ships in it. And then we go on to the various requisitioning programs in which the government decrees that it is a national emergency of some sort. Thus, if we cannot pick up ships in the voluntary market they will be made available to us.

The major problem is the positioning of ships. Aircraft can be repositioned; in a day you can change the whole complexion of the aircraft fleet. You cannot do that with ships but you have many more ships. So sometimes we have to step through various programs before we can obtain a ship.

We match the requirements to the assets. At this point, we are not looking for a schedule or even a definite fleet sizing. We are looking for a date to complete the lift. Whatever it is, we give a number, and then we give any assumptions, such as "we are going to need to break out the RRF" or "that is assuming we can get six voluntary charters."

KASKIN: Under the crisis action system that they are currently trying to develop, would not there be in that notional lift a little bit more information, almost similar to the detailed planning information, so that you can translate the cargo lift into the type of cargo and when they want it to arrive, so that you have dates involved?

BALLOU: At this point today, there are no dates. There are no RDDs in it. This is where you get your chance to say what you can do. You are talking about something they are looking at for the future.

As an example, consider the airlift command (MAC). They have a desk top calculator that does this. They assume that everything is going to lift out of the middle of the country and go on to Europe, Japan, etc. That is the kind of calculation they do. It is very rough. So at this point, we are not looking for something sophisticated.

JAIKUMAR: Are you saying that the objective, in some sense, is to minimize the maximum time to complete the lift?

DOUGLAS: With a given fleet.

BALLOU: That is a good point. We might want to come up with three dates, one just using MSC ships, one using MSC ships and the RRF, and one if we go to requisitioning.

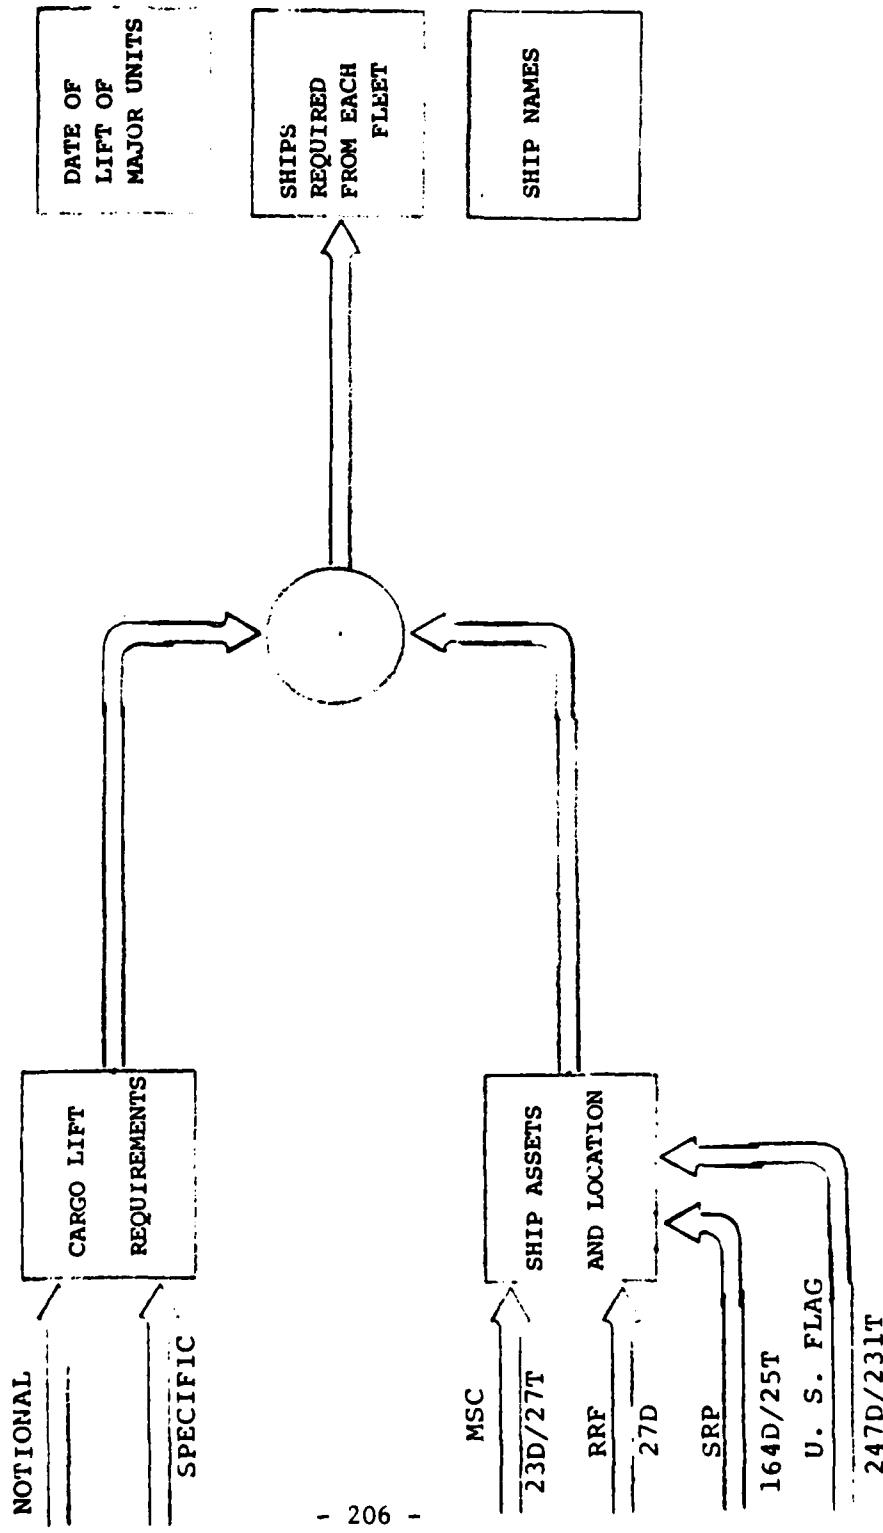
KASKIN: I have two comments on ship asset availability. The timing of the availability of those ships is different. For example, a ship coming out of the national defense reserve fleet will not be available for 30 to 60 days. Additionally, if it is a European war, we have NATO vessels that we can add to the fleet, and that is another several hundred vessels, so the problem would be larger and more complex.

BALLOU: In Figure 9.2 is shown the next phase, but it is still part of what I call the force sizing phase. We may be starting to get more specific information now. In a sense that creates a problem because then you have to sort out the information in terms of the notional lift and what is now more specific.

But that is part of the input, not the actual solving of the problem. The fleet is the same as before. You can see what we are required to give is a little more information, a little structure, so to speak, to the delivery points, but still not the details. Here we are, starting to determine the numbers of ships. How much of the RRF do we need to break out? How much should we ask for with respect to the SRP? Basically, that will answer our requirements at this stage.

FIGURE 9.2

FORCE SIZING DECISIONS III



But, obviously, we have to start getting down to the specific ships. We have to know the names of the ships because of the location problem. This is a change in Figure 9.2; I did not have location in Figure 9.1. Usually you indicate it if you do know where a ship is, but here you have to start really looking at the details of the ships and making the RDDs something that you can meet.

Let me mention one of the problems that we see in the way we do it now. We looked at some of the deliberate planning models that were mentioned yesterday. There is the SEACOP model. There is the SEASTRAT development. We presently take a look at a plan we have and modify it as necessary by pencil, crossing out, adding in, etc. Obviously, that is not very easy to do if you have a large number of ships.

There is the port call problem. Sometimes the model will have us going in for 10 tons of cargo. And then there is the multi-port problem, six or eight calls on each side of the Atlantic. It just does not seem to make sense. And yet, on top of that, the model sails ships that are only half full.

DOENGES: Would you clarify something? Are you saying that you do additional SEACOP runs yourself or are you pulling previously developed and delivered planning off the shelf, seeing if it is appropriate, and marking it up for your purposes?

BALLOU: Our only present capability is to do the latter. We do not have any capability to reflow. That is in contrast to what the air people do. They have reflow capability in this area.

In the phase previous to this one, if we get some idea of the amount of tonnage, then to be able to get the number of ships we divide by something like 10,000 and that tells us the number of ships needed. And then we look and see what is available on the Atlantic coast, and we have the solution. However, this needs more sophistication.

BISHOP: That must be a very rough approach.

BALLOU: Yes indeed.

BISHOP: Ten tons of tanks is a little different shipping requirement than 10 tons of jeeps. You are probably going to go by metric for most of your cargo, rather than by weight.

BALLOU: Yes. When I said 10 tons I was talking about 10 measurement tons, using the volume ton.

BISHOP: But is that your real constraint?

BALLOU: Yes.

POLLOCK: Vehicles would be square feet.

BISHOP: And troops would be too, I would think.

KASKIN: Right now we do not have the detailed information early in an exercise. The Joint Deployment Agency is trying to build up a capability so that it would be available quickly, so that if somebody said they wanted a particular force you could get the movement requirements generated immediately, and then it would be available for processing. Today, that is just not possible.

BALLOU: There was an effort at Texas A&M back in the early 1970s to look at our tanker scheduling business. This was for peacetime, normal operations. As best I can tell, the conclusion, at least from the operator's point of view, was that it did not do any better than what they were doing manually, and so it was not continued.

The closest thing we have now is really more of a bookkeeping operation. We have a program we call CALSTAT, Cargo Location and Status, which deals with ship location and status, as well as with cargo. It basically keeps everyone playing from the same sheet of music, i.e., everybody knows the latest versions of the ship schedules.

Then there is another program that we have, one we are really using more in peacetime than in wartime, if we ever get down to this detail. It is Proforma, and it analyzes individual ship

voyages. You put in the ship characteristics, the cargo characteristics, and then you play with the ports. What it amounts to is that you get out what you put in. If you want to run sensitivity analyses, you just make more runs and see what makes sense as far as port ordering goes, etc.

GREER: It has a financial aspect.

BALLOU: Yes, it yields cost information.

KASKIN: Those programs are not sized to deal with wartime.

BALLOU: No, they are single-ship programs.

Now, with the addition of the actual scheduling shown in Figure 9.3 you have the three phases required in our operation. The only other information I have here, if it would help, is some idea of the amounts of cargo and the numbers of ships that we have dealt with in exercises. I think Captain Scott mentioned yesterday that about once every two months there is an exercise or a real world situation where we get asked "what if" questions. And these certainly do not deal in hundreds of ships or millions of tons of cargo. They are usually small scale things. But, certainly, we need to provide better information for the large problems because you do not have statistics working on your side to come out with the final numbers.

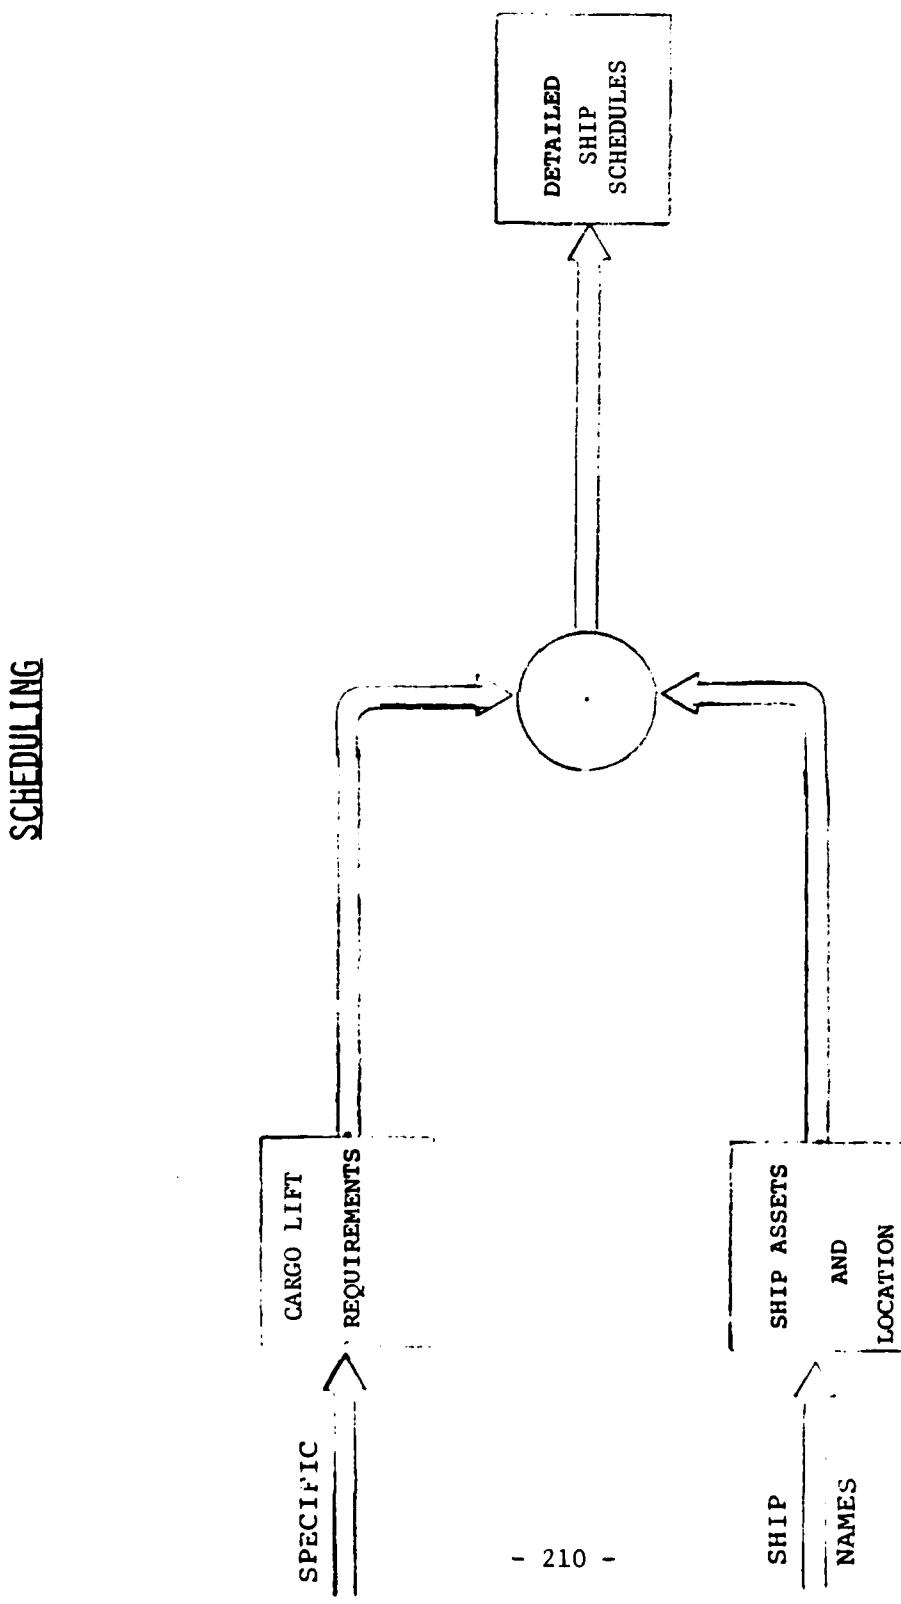
Now, I would like to ask you about the similarities or the differences that you see between some of these designs and those that have benefited the commercial world. Are there ones that could directly transfer over to us? Are there certain areas that we should look at?

RONEN: Basically, I see the problems as quite similar. These three pages here (see Figures 9.4, 9.5, and 9.6, respectively) are similar to long-term, medium-term, and short-term planning.

BALLOU: Exactly.

RONEN: You also have problems in short-term planning where you have to assign ships to cargoes and specific cargoes to specific ships, which

FIGURE 9.3



CRISIS MANAGEMENT
COURSE OF ACTION DEVELOPMENT

PURPOSE: PROVIDE QUICK ESTIMATE OF SEALIFT CAPABILITY
DURING A CRISIS

OBJECTIVE: TO IDENTIFY CLOSURE DATES FOR MULTIPLE COURSES
OF ACTION USING NOTIONAL SETS OF REQUIREMENTS,
THE AVAILABLE POOLS OF SHIPS AND OPTIMIZE TO
EARLIEST COMPLETION DATE

SCOPE: SHIP REQUIRED (1-600)
SHIP AVAILABLE (40-1200)
UP TO 150 POE, 80 POD'S

TIME HORIZON: 14 TO 180 DAYS

FREQUENCY: AT CRISIS TIME AND 3 OR 4 EXERCISES/YEAR

TIME RESPONSE: 1 HOUR

CONSTRAINTS: PORT CONSTRAINT, CARGO/SHIP CHARACTERISTIC AND
AVAILABILITY CONSTRAINTS

CURRENT
CAPABILITIES: WITHIN STATE OF ART

FIGURE 9.4

CRISIS MANAGEMENT
EXECUTION PLANNING

PURPOSE: PROVIDE DETAILS ON SELECTED COA

OBJECTIVE:

- o TO IDENTIFY SHIPPING ASSETS NEEDED FOR SELECTED COA
- o PROVIDE CLOSURE ESTIMATE ON VARIOUS AGGREGATIONS OF CARGO
- o OPTIMIZE TO LEAST COST

SCOPE: SAME AS COA

TIME HORIZON: 14 TO 180 DAYS

FREQUENCY: AT CRISIS TIMES; 3 OR 4 EXERCISES/YEAR

RESPONSE: 4 HOURS

CONSTRAINTS: PORT CONSTRAINTS
CARGO/SHIP CONSTRAINTS
CARGO UNIT COHESION

CURRENT

CAPABILITIES: WITHIN STATE OF ART

FIGURE 9.5

CRISIS MANAGEMENT
EXECUTION

PURPOSE: MANAGE THE SHIPPING ASSETS OF MSC DURING A CRISIS

OBJECTIVE: PROVIDE DETAILED SCHEDULES WHICH WILL MATCH SHIPS WITH CARGO SUCH THAT REQUIRED DELIVERY DATES ARE MET AT LEAST COST

SCOPE: SHIP REQUIRED (1-600)
SHIP AVAILABLE (40-1200)
POE (1-150) POD (1-80)

TIME HORIZON: 2 TO 5 WEEKS

FREQUENCY: AT CRISIS TIME 3 OR 4 EXERCISES/YEAR

RESPONSE: 2 HOURS

CONSTRAINTS: PORT CONSTRAINTS
CARGO/SHIP CONSTRAINTS
CARGO UNIT COHESION

FIGURE 9.6

is the problem that everybody else has. There is definitely a problem and it has not yet been solved properly. As I say, the major direction now is in man-machine systems, interfaces with on-line systems.

The major question I would like to ask you is this: suppose you are presented with two different schedules or loading plans. How would you evaluate them? On what basis?

BALLOU: Are you talking in terms of optimization?

RONEN: I was not talking about optimization. I just want your criteria for evaluating schedules.

BALLOU: Our major objective?

RONEN: Yes, your major effectiveness or objective function.

BALLOU: One of the things that I think has driven this other model that was mentioned yesterday is this idea of the delivery date of the cargo. The penalty for not delivering it can go anywhere up to infinity where you just do not allow that to happen, and that, of course, drives the number of ships that you use to infinity as well. I think that is something that we have not really addressed properly. RDD certainly is a major consideration.

RONEN: Does that drive the whole system?

BALLOU: You have to get the cargo there on time. And then another point that I did not mention is that we cannot refuse cargo. We cannot say, for example, that it just does not make sense to go to Port B to pick up cargo. That is not an option. We have to do it.

And so the thing we are going for is lifting all cargo that comes in as a requirement and delivering it on time.

KASKIN: With the fewest number of ships.

BALLOU: Yes, with the fewest number of ships. Every time we step to a different level, as you saw there in the fleets, a different level of national interest, and additional justification to do it is required. MARAD, for instance, makes an analysis of the impact on the economy. Even if you are talking about a full-scale war, it may be more important to bring in iron or aluminum, or particularly oil. You cannot disrupt certain things just because somebody wants something.

So we work with MARAD in these areas with respect to the numbers of ships. There is a limit to the number of ships that we would have access to.

POLLOCK: It is even more complicated than that. When we allocate ships to you, what we do is set aside our requirements for domestic needs. You mentioned the NATO ships, for instance, as one of the resources that is available to you. There is a pool of 600 ships, of which you can call upon 400 in a crisis.

Even there, the situation is such that in the allocation of those ships you have to be careful because you cannot allocate ships of only one country because you would have the Director General of the CSG and NATO management complaining through the White House and the State Department stopping you from tearing up the merchant marine of one country versus another country. So you have a different problem. You have to pick ships from different countries; it is really a very complicated problem. It has a political overtone that is very real and in the last exercise we had, Greece pulled out of NATO.

It is not just a straightforward analytical problem. It is a problem that requires a lot of judgment. So I think you are going to need an on-line system to bring in a lot of the input to solve this problem.

GREER: It seems to me if you cannot meet your requirements and you are thinking about your SRP ships, from a MARAD point of view, maybe what you should do is to have all of those carriers and all of the American companies come in first and tell you what is available and where it is.

POLLOCK: That is right. There are availability models which will tell you about any port you want. We use a generalized scheme, e.g., east coast United States would be some port like Norfolk. We can tell you what ships will be available to you very quickly. In fact, you could almost have that data any time.

BISHOP: As far as I understand that system, though, they track ships. They do not know the cargo condition of a ship. Having a tanker that is loaded may not do you any good.

BALLOU: I think another aspect about either the force sizing or the scheduling is that it is a time-phased problem. Also, it is hard to say what the initial conditions are; it is not like all the trucks are in the garage when you start out in the morning.

I do not know whether that is a problem for scheduling algorithms; whether that introduces additional complexities or makes things simpler. But the initial condition of the assets is a bit of a complicating factor.

JAIKUMAR: For the first phase, which you give in Figure 9.4, you have as clean an objective as you can possibly get for any model. Given what you have here, and talking from an optimization point of view, I think that this problem can be optimized; if not optimized, you can have a heuristic which will tell you how far from optimality you really are.

BALLOU: Well, let me back up a minute on that. Let us take a look at that objective.

There was some discussion yesterday about how we are dealing with cost. Maybe that is a valid point. I have not thought of anything better than earliest completion date, and that is why I ended up with that.

GREER: I can partly see why, perhaps. Suppose you are going to Port A on the east coast and you need eight ships there. And you are going to go to two or three locations overseas. You plan that you are going to put this unit on that ship and a second unit on another ship. In reality, as those ships arrive in port A, if there is any urgency to the situation, you are probably going to put whatever cargo is available on the first ship that arrives, rather than sit there and wait for two or three days. But, is that ship as fast as the ship that you originally planned on using?

You have to have an iteration there, a control.

KASKIN: I see your point. But how do you plan for that? And what should I be doing given what I have, or what I think I have, which is what I think Figure 9.4 is about. Given this statement in Figure 9.4, I do not see why one should not at least attempt to look at it from an optimization perspective.

FOARD: That is a good statement. Does anybody in the community take exception to that? The scenario that this occurs under is that a crisis has occurred somewhere in the world and the military has been alerted to respond to it. The theater commander out there has proposed five or six different courses of action through the national command authority here. Those courses of action are put out to MSC, to MAC and to various other commands to evaluate and recommend an action. That is why we are identifying multiple courses of action.

A response would go back to the national command authority saying that we cannot support those two. We can support these two. And a decision would be made to implement one of those. So that is the kind of analysis we are supporting here in the first phase.

The next thing that will occur is that we will implement one and then we will go to matching a ship to a cargo to a port, the scheduling problem. But the first phase is an overall gross feasibility exercise of a course of action, given a group of forces that are going to go some place, etc.

BALLOU: The biggest problem, really, is identifying the groups of forces.

FOARD: If you cannot identify the groups of forces you cannot do the thing to begin with. So you have to assume that you are going to be able to get information about the forces in order to do the feasibility study.

ENGLISH: May I go back to the question of least cost? It seems to me that if you deal with least cost in some simpler way rather than dollar cost, say minimum number of assets or minimum number of ships, without any fancy regard to the dollars and cents, that might be better. In the defense establishment there is an expression "cheap dollars and expensive dollars," and once you are involved in a war, the dollars are cheap. I see problems if you are really going to have an accountant in there counting the pennies. It seems to me that least cost should have to do with some gross measure of assets--economy of means (e.g. ship assets).

BALLOU: Cost is a vague term. Whether you want to attribute it to dollars or assets, I think at this stage it has to be kind of an output because we are getting questions that way now, e.g., "how much will that cost?" It may not be used as the trade-off. Maybe we should not optimize against the cost.

ENGLISH: It is possible, Joe, that the two are fairly close. And I do not think that you should agonize over the dollars too much in the beginning. And also, it seems to me that you are given a hierarchy of availability because of the fact that there are political difficulties with certain ships, they get bumped down in your prior planning arrangements so that they stand low on the list. And as you work your way down through the list, being as economical in the use of your assets as you can be, that could conceivably, at least in a zero configuration, resolve some of these other judgmental problems that are deeper.

KASKIN: I think it needs to be emphasized in this course of action development that we do not have detailed information. And that is

one of the reasons why the problem is as fuzzy as it is. But what we probably would do if we could get this capability would be to develop various closure dates for various scenarios that are specified. Some of those scenarios are based on what fleets of ships we assume are available.

Another aspect of it that may or may not apply is whether we have convoying or not. If we assume independent sailing versus convoying, it is going to make a major difference as to when those ships come back. Both of these force-sizing models have to assume multiple voyages. I do not think that the actual schedule would go beyond one voyage because during a real war you just do not have enough information. You are really not going to be able to predict when a ship will actually come back. But, if you can at least assign the initial cargoes to the initial set of assets, you have solved a major part of your problem.

But, in these other two aspects of course of action development, the other force-sizing aspects of it, you are going to have to work along the horizon and it makes the problem difficult when treating these other factors, such as attrition. There has to be something in the model to account for them.

HALL: There is something that bothers me about the objective in Figure 9.5. From our point of view at the Joint Deployment Agency, if we were actually getting ready to engage in a conflict, I do not think that we would be so concerned about the efficiency as we would be with the capability that you can project.

Our figure of merit would be how much combat power you could project. Now that is a very difficult thing to quantify and it would have to be done subjectively today because we do not have any good metrics or models to measure that.

And to answer the question that was asked a while ago, if I were presented with two schedules and you asked me which I liked best, and this is assuming that they both met what we perceived to be the required delivery dates, I would have to give them to the supported commander and ask do you have a preference? What he would probably do is look

at them and see if the stream of combat forces arriving was significantly different, and if one would give him a better combat advantage than the other. He would probably pick the one that gave him the best combat power.

DOUGLAS: He is going to get the stream. If it follows the C plus 5, 10, 15 timeframe, he is going to get that.

HALL: Okay, but they are going to be different. The flow is going to be different. The sequencing will be different.

DOENGES: Let us come back to their first phase here, where I believe that Joe Ballou has said that they do not have delivery requirements by date. So if you are looking at putting together what will help you on the first phase, how are you going to break ties? You need a value scheme, e.g., it is more important to load one piece of additional cargo on this notional ship than it is a different piece.

ENGLISH: The fewest number of ships versus the time; that is the objective.

BALLOU: In some ways I regret putting down that objective, and I know it is different from what I said when I was talking.

HALL: Let me finish and I will be brief. I think what bothers me about the objective is that, as I read it, we are subjugating the closure date to the attainment of least cost. And if I were the commander, that is not what I would be worried about. I would want you to minimize the closure date, even if it would cost me more ships.

ENGLISH: Let me suggest what I think is perhaps part of the solution. You people in JDA look at the priorities and what the military realities are, and you try to get these built into the system as prior decisions. Essentially, then, the optimization takes as given that you must deliver the important stuff first. That leads me to bring up an issue that came up yesterday and, to my view, just lay there.

Do you remember when the question was asked--what is the measure of priority? We were all able to tell ourselves that RDD was the measure of priority. Then someone else and I both asked if there are other additional measures that further define or further explain the priority?

I ask the question bluntly; if you have two bunches of stuff and they both have the same RDD, are they viewed as equally important? Now that may be true. But I wonder, is there a degree of subtlety that the commanders, the military guys who are going to do the fighting, could further define, so that you could better understand which was more important between two bunches of stuff.

I guess what I am suggesting is that it would be nice if you had an a priori set of values built onto the stuff to be moved, all the way through. Of course you may want to fine-tune them because the war changes. But it seems to me that in these first cuts, that that would take care of a fair bit of the difficulty and then the optimization could, to some extent, go on its merry way because the values would already be built into the model.

HALL: I think the difference, though, is whether you are trying to minimize the delivery date or whether you are trying to minimize the number of assets it takes to accomplish the total deployment. I think it is different.

ENGLISH: In terms of assets, I think it depends on whether or not you are ship-rich. Does it somehow turn out that you begin the whole exercise with all the ships you are ever going to need? Yes? No? If it is yes, it is possible that there are all kinds of things you can do to really beat out the enemy in a military sense. If you are ship-lean, however, then you really do have to husband your assets because of sinkings and later phases of the war.

FOARD: We have two different problems here. Suppose we build what is known as a TPFDD and that it is used as a movement requirement supportive of a particular operations plan; it is against this that the deliberate planning phase is run which is presently being treated by Discussion Group 4. We are the crisis action response group here. We may pull

that TPFDD off the shelf if it is the right TPFDD, the right operations plan to execute. It has short tons and measurement tons of resupply, ammunition, army support troops, army support equipment, etc. What really happens when it comes down to who is going to select what goes on what ship? It is going to be the guy over there, because we are not going to take all of the food and ammunition that we have sitting around in warehouses and depots and move it down to the pier and start loading it on a ship.

We are going to do that for about the first ten days just to give that theater commander over there some support. From that point on, he is going through normal supply requisitioning procedures to move that stock to him and he is, after about the first ten days, going to be telling us what his priority items are by what he orders. He will have regular routine stuff moving, but he will have top priority stuff too, and that is where we will get the relative values.

ENGLISH: I know from Woody Hall at JDA that the military are working these problems. Is it possible that there is another indicator rather than RDD and that there should, perhaps, be a system using it?

HALL: Yesterday I did not know how much to say. I know that in the planning system, the JOPS itself, although I am not a JOPS expert, there is a priority indicator.

KASKIN: It is priority on that day of that particular RDD which the CINC can specify. But it is not on one RDD date, day 14 compared to day 15.

HALL: That is true. It is within one day.

FOARD: The guy who builds the lift does have the capability to assign a priority to do the different things with the same RDD.

BALLOU: I think we are getting off into another problem, the definition of the input of the cargo. I would like to go back to focus on the scheduling and modeling aspects of the force sizing. And I would like to take a look at some of these other phases as well.

COPPINS: I have two quick comments. The suggestion was made that it depends on whether you are ship-lean or not. You may not know if you are ship-lean until you know what kind of closure dates you are talking about. If you have 180 days to do something, you may be able to do it with a very small fleet. If you have to do it in ten days, you may need a large fleet.

This is a multidimensional problem in terms of optimization. If you are going to minimize assets, you really need to look at it parametrically, by closure date. That is one way to do it; treat it as a clear multi-criteria problem. Then there is not an optimal solution. There may be a minimum asset requirement given a set of closure dates, and then it is going to be up to somebody to decide on the trade-off.

I think that is what you are really looking for; it is not this set of closure dates, or what are the minimum assets needed. It is, "Okay, here is what I can do on this date. Here is what I can do on that date. What do you want to do?" And if there is one course of action that you cannot possibly achieve, then it is one that you can discard. But my guess is that you will probably have possibilities for every course of action. And then somebody is going to have to make a management decision.

Getting into the other problem; it is the same problem that Bruce Bishop referred to. It is the same thing that I saw at Reynolds Metals. Your first two phases deal with a long-term or, at best, a medium-term situation. Whatever you optimize it with, whether you use an LP or a more detailed mixed integer programming model, how do you translate the results into the detailed scheduling. The initial results may identify something that you can do with a certain number of ships, but when the reality comes up that you do not really have this ship in port and it is stuck out in the middle of the ocean, and this one is full of ballast and that one is in drydock for an extra week, how close is what you identified before going to match what you can actually do with a detailed schedule? That is something that I am not sure anybody has resolved yet.

DOUGLAS: But you can find out. You can exercise both systems, because this is all ahead of time, anyway. You can set it up so that you get your first modeling results and see that it looks as though you can do it with this many ships. Then you plug in the scheduling model, the more detailed stuff, and depending on where those ships are at this point in time, where they are now, suppose it looks all right. You run it again a month from now when the ships are in different places, and, golly, it does not work out.

If you want to get real fine, you run it a thousand times and see what the probability is that you are going to be able to meet your requirements. You can play with it.

BISHOP: Given the constraint that you indicated, that you cannot miss the RDDs, I do not think that you can separate the problem. I do not think you can take a look at the problem as a pool of tons or as a pool of ships and then run an LP to find out what you are doing over time and then commit yourself to it without looking at the short-term problem.

BALLOU: Are you saying you have to do scheduling right from the beginning, right up front?

BISHOP: Given the fact that you have to commit yourself to dates and absolutely hold to them, I do not think that you can break them apart.

POLLOCK: I have a suggestion that may help out here. There are certain assets that are in more critical supply. For example, among your shipping assets there are not many ROROs available. Those are the ones that you might want to monitor. With respect to container ships, there are probably a lot of them available and, therefore, for any problem that comes up, the probability of container ships being available is high. The same thing is likely true with freighters.

So perhaps, what you should do is determine the critical assets and then I think you should look at where they are and continuously monitor them. That way, I think that with almost anything you come up with in your plan, there will be enough resources available to have a high probability of success in the plan.

KASKIN: Is there a consensus that you need to use a scheduling algorithm, even if the course of action is over with, in order to make those determinations?

DOUGLAS: No, I do not think so.

RONEN: I do not agree. You cannot schedule multiple voyages. It is impractical to schedule multiple voyages. The reason is that if you have attrition, you have uncertainties and you do not know what ships will be available for second voyages.

BALLOU: I think when you are scheduling you do not assume that a ship is going to be lost.

RONEN: You do not have to assume it. But you do know that some of them will be lost, and some will be delayed.

BALLOU: Look at Figure 9.6 on execution. Perhaps I am optimistic, but I think I put down two to five weeks; that is the time interval I am interested in. Basically, I am thinking of a one-voyage situation. You step it ahead each day and you look ahead two to five weeks.

RONEN: For the long-range planning problems you can try to work with averages. But a day's average would have to be inflated to take into account the losses and delays.

GREER: Getting back to the cost, it can be a gross term and it should be, but one of the most embarrassing questions to try to answer is, "What is it going to cost?" It involves perhaps 30 ships or 90 ships, and somebody over in the JCS asks what is it going to cost to do it?

FOARD: That is my point also. They sit around with nothing to do in an exercise and wonder what it is going to cost. So they ask that question. We are talking here about developing a system that is going to respond to an actual crisis. Nobody is going to ask about cost.

GREER: When they started up the ammunition pipeline to Vietnam, the first questions that were asked were about cost. What is it going to

cost if we have to go around South America? What is it going to cost if we go through the canal? What is it going to cost if we go via the east coast or west coast? That came from MSC. It took us over three months to answer the questions.

All I am saying is that you ought to tie a gross cost figure to this thing because somebody will ask. Everyone gets involved, all the agencies of the government, the Congress; I would not overwork the problem, but I would at least do it in gross terms so there is an interim answer.

BALLOU: I would like to get your impressions about what I have said here in Figure 9.6 on execution. Do you still have comments on force sizing, working out some of the details? Maybe we do not have the objective statement just right, but it is at least something that looks like it may be reasonable to look at and do. How about the scheduling, what I call execution?

DOENGES: I have a question for you, Joe. Someone said something to the effect that there would always be a way to meet the required schedules. Do you agree with that? I believe he said something to the effect that unless it were a really strange situation, there would always be at least one course of action that could meet the objective.

ENGLISH: We may have been talking about the ship-rich situation.

COPPINS: That was back in the long-range planning mode where if you fool with the closure date, and you move it out far enough, then the answer is yes.

DOENGES: It is certainly possible that you could run into situations where, on a first pass, the problem is not feasible and you have to change things around in order to find a way to meet the schedule.

GREER: The real problem is probably whether you can deploy fast enough. If you cannot do it fast enough, maybe you are not going to do it at all. You have lost the objective.

DOENGES: I would just like to suggest that it is certainly worth exploring optimization methods to attack your execution task because, in those situations where the problem is not feasible, the optimization techniques give you a lot of valuable information as to what to do in order to attain feasibility. With something like a simulation approach, however, that information might be more obscure.

BALLOU: Are there practical scheduling techniques that would work with the kind of problem we are talking about?

RONEN: Is this short-term scheduling?

BALLOU: Yes, we are talking about the short-term problem. Roughly speaking, the horizon is the voyage.

RONEN: That is an assignment problem.

JAIKUMAR: The size of the problem that you are talking about is eminently feasible.

RONEN: It is an assignment problem, cargo to ships.

JAIKUMAR: You can even do multiple voyages, if you wish, in the time horizon that you are talking about. It does not necessarily mean just one voyage.

BALLOU: It is not a two-port problem.

JAIKUMAR: The point is that you can generate a set of feasible routes similar in structure to what I talked about yesterday. The routes will contain at most five ports. I am just giving you a rough number when I say five ports. And you can have multiple voyages. The problem size you are looking at suggests that, at most, you are going to look at about 100,000 possible combinations.

Discussion Group 4 - Contingency
Planning and Scheduling

Leader: Lawrence A. Gallagher
Military Sealift Command
Department of the Navy
Washington, DC

GALLAGHER: I am not happy with the results of SEACOP, and that is one of the reasons we are here. We are trying to fix that problem, to come up with a different way of doing business. It would not be a military presentation if I did not have a handout, and this one (see Figure 10.1) summarizes some of the issues to be discussed. Just to get your thoughts around to the scope of the problem that we are faced with in the military, I would like to review the problem and take a look at our requirements and some of the discussion topics (see Figure 10.2)

A primary mission of the Military Sealift Command is the strategic planning of sealift operations for contingency and mobilization situations. The key requisite is the capability to move military forces and supplies to the required location within the required time during a period of potential or actual conflict. For this reason, MSC is currently developing a comprehensive methodology, SEA STRAT, to perform the scheduling of MSC shipping transportation resources and thereby evaluate the feasibility of meeting mobilization requirements.

You heard Colonel Wilmes when he spoke yesterday. SEA STRAT is more than just a scheduler. SEA STRAT will be a system that is going to help us with our long-term planning. I think we have now a total of six separate modules that we have identified, six different areas that will require modeling.

DISCUSSION GROUP 4 CONTINGENCY PLANNING AND SCHEDULING

A primary mission of the Military Sealift Command is the strategic planning of sealift operations for contingency and mobilization situations. The key requisite is the capability to move military forces and supplies to the required location within the required time during a period of potential or actual conflict. For this reason, MSC has initiated the development of a comprehensive methodology, SEASTRAT, to perform the scheduling of MSC shipping transportation resources and thereby evaluate the feasibility of meeting mobilization requirements.

The objective of the SEASTRAT model is to schedule ships and cargo transport so as to deliver cargo to the required port of debarkation by the proper time. The ship routing and scheduling must be performed consistent with initial cargo and ship locations, port capabilities, cargo and ship types, and efficient ship utilization so as to best achieve delivery requirements.

The SEASTRAT problem size can range from 5,000 up to 15,000 cargo units to be delivered during any one analysis. The number of available ships can range up to 1,000 ships that would be under MSC control in any given scenario. Ports could include 150 ports of embarkation and 80 ports of debarkation. The time horizon for analysis is usually 180 days, which means that multiple voyages are possible for each ship.

Discussion Group 4 will address several generic methodology issues for SEASTRAT development. For example, one key issue is level of detail, which can range from gross feasibility to a detailed, executable schedule. Gross feasibility answers such questions as: are there enough ships to lift the cargo in the time required (irrespective of details such as port throughput, convoying, etc.)? An executable schedule, on the other hand, is so detailed that the cargo and ship movements could be implemented in the "real world".

A second issue is the actual scheduling algorithm to be used in SEASTRAT. Heuristic methods use reasonable logic and decision rules to schedule ships and cargos - the problem comes in defining what is "reasonable". An optimization approach uses mathematical techniques to determine the "best" solution; however such techniques may not be computationally feasible for large problems. A hybrid approach could combine both heuristic and optimization methods.

An additional consideration is the need for modeling approaches which improve computational tractability. Aggregation can be used to combine small elements such as individual cargo, ships, or ports into a larger element, which can then be treated as a single unit without altering the basic features of the problem. Partitioning or decomposition techniques can be used to structure a high-level, large model into smaller submodels which can then be solved nearly independently. For example, individual ports could be partitioned into port regions. Approximation can be used to model complicated relationships in simpler terms, with subsequent improved approximations made as the solution is obtained.

SEASTRAT also has a number of special features which could be addressed in the Discussion Group. These features include the need for convoying, the statistical nature of ship locations at the starting point of a mobilization requirement, and the need for interface with land transportation facilities.

FIGURE 10.1

DISCUSSION GROUP AGENDA
CONTINGENCY PLANNING AND SCHEDULING

- REVIEW OF SEA STRAT SCHEDULING PROBLEM
- SUMMARY OF MSC REQUIREMENTS
- DISCUSSION TOPICS
 - APPLICABILITY OF PREVIOUS SCHEDULING WORK
 - HEURISTIC VERSUS OPTIMIZATION APPROACHES
 - APPLICABILITY OF SPECIFIC SCHEDULING TECHNIQUES
 - METHODS TO IMPROVE COMPUTATIONAL TRACTABILITY
 - USER INTERACTION
 - OTHER

FIGURE 10.2

FIGURE 10.3

SEASTRAT MODELING OBJECTIVE

- SCHEDULE SHIPS AND CARGO TRANSPORT SO AS TO DELIVER CARGO TO THE REQUIRED PORT WITHIN THE REQUIRED TIME, CONSISTENT WITH
 - SHIP LOCATIONS
 - CARGO LOCATIONS
 - PORT CAPABILITIES
 - CARGO AND SHIP TYPES
 - EFFICIENCY SHIP UTILIZATION
 - DELIVERY REQUIREMENTS

The objective of the SEASTRAT model is to schedule ships and cargo transport so as to deliver cargo to the required port of debarkation by the proper time. The ship routing and scheduling must be performed consistent with initial cargo and ship locations, port capabilities, cargo and ship types, and efficient ship utilization so as to best achieve delivery requirements.

An overview of the SEASTRAT transportation feasibility model is shown in Figure 10.4. Major input categories include:

- cargo data
- ports of embarkation and cargo availability dates
- ports of debarkation and arrival dates
- ship data such as location, speed, type, etc., and
- shipping strategy factors such as preferred sea lanes or the need for convoying.

The SEASTRAT model will then perform ship routing, scheduling, and cargo delivery with MSC interaction to specify additional constraints or schedule requirements. A key output of SEASTRAT is feasibility--are the cargo delivery requirements met? The feasibility evaluation should be available as both a gross, preliminary feasibility estimate and as a detailed ship-by-ship and cargo-by-cargo analysis of shortfalls. In addition, for detailed analysis, SEASTRAT should provide comprehensive cargo movement schedules, ship routing and scheduling, and identification of slippages and other problem areas.

Our inputs are modular and so are our outputs. We now have a model that is called SEACOP, but we are going to have SEASTRAT modeling. We do have all of these inputs to our system at this time. We do have MSC/User interaction, but it is very limited. When the system was developed it was a front-load system. You put the data in, and you waited to get the results. You could not do anything in between. I found a few ways to change the data base around so that I can make the logic in the scheduler do anything I want by changing the data base. I now have, instead of just one data base, a series of data bases that I use to come up with some predetermined results. I know that it is not right, but that happens to be the way it works right now.

SEASTRAT OVERVIEW
SHIPPING TRANSPORTATION FEASIBILITY ANALYSIS

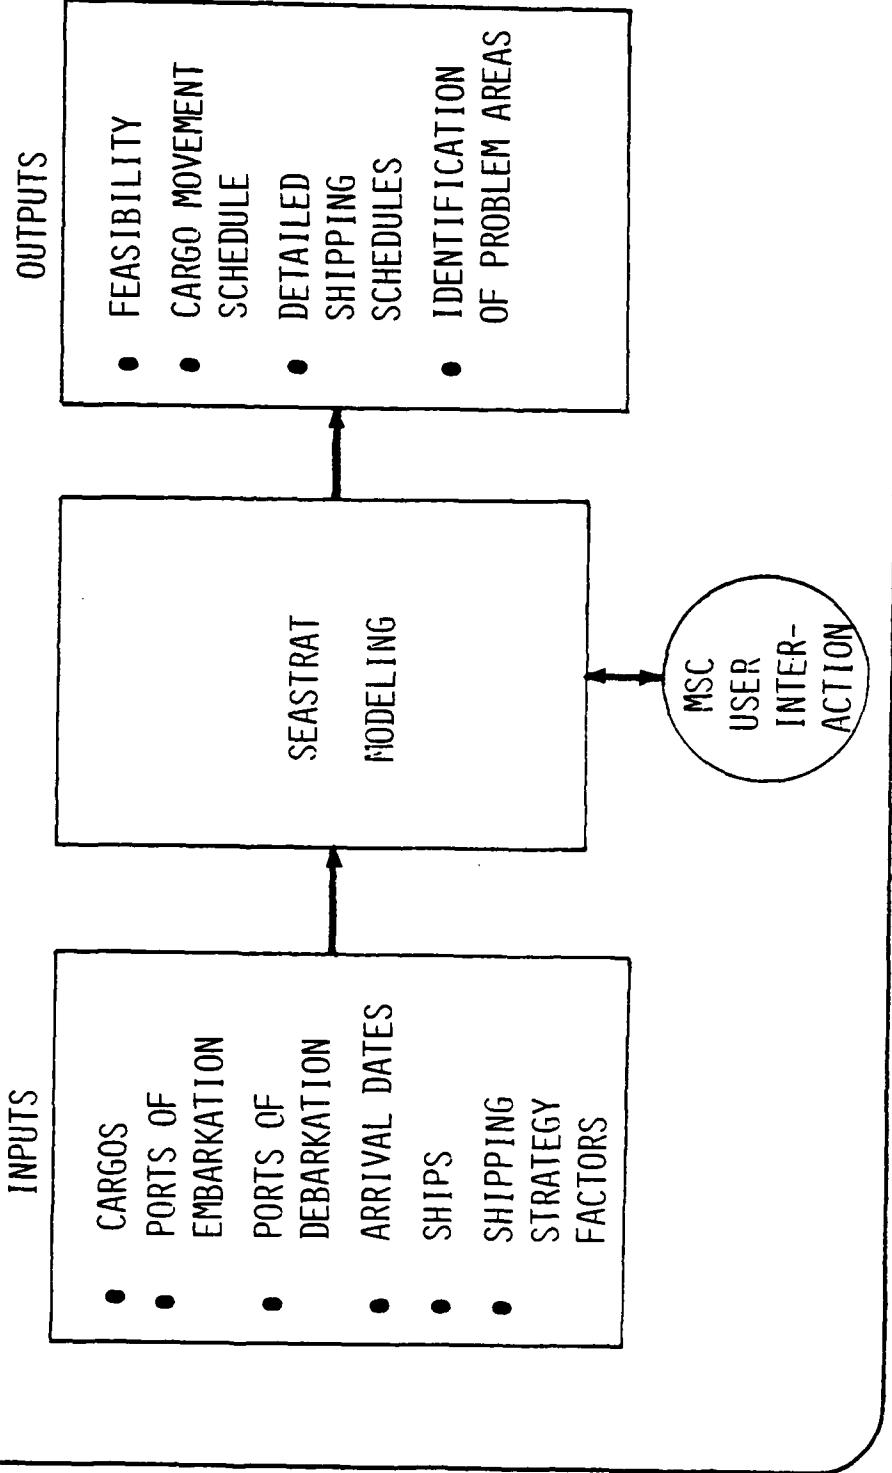


FIGURE 10.4

SEA STRAT PROBLEM SIZE

- 5,000 - 15,000 INDIVIDUAL CARGO UNITS TO BE SHIPPED
- UP TO 1,000 SHIPS
- 150 PORTS OF EMBARKATION
- 80 PORTS OF DEBARKEATION
- 180 DAYS

FIGURE 10.5

INPUT DATA FOR SEA STRAT
TRANSPORTATION FEASIBILITY ANALYSIS

- DESCRIPTION OF CARGO UNITS
 - CARGO CHARACTERISTICS (TYPE, AMOUNT, DESCRIPTION)
 - PORT OF EMBARKATION (POE)
 - PORT OF DEBARKATION (POD)
 - EARLIEST AVAILABILITY AND REQUIRED ARRIVAL DATE
- DESCRIPTION OF PORTS
 - PORT CHARACTERISTICS (E.G., THROUGHPUT, DRAFT, EQUIPMENT)
 - LIST OF PORTS OF EMBARKATION
 - LIST OF PORTS OF DEBARKATION
- SHIP DATA
 - SHIP CHARACTERISTICS (SIZE, SPEED, LOADING, CARGO CAPABILITY, ETC.)
 - SHIP AVAILABILITY AND LOCATION
- SHIP ROUTING DATA
 - SHIPPING LANES
 - OCEAN CLEARANCE
- SHIPPING STRATEGY FACTORS
 - CONVOYING
 - SHIP OR PORT ATTRITION
 - CARGO UNIT COHESION
 - RESUPPLY REQUIREMENTS

I have already decided the way things should come out, and then I feed my data base to make them come out that way.

Figure 10.5 deals with the problem size. The SEASTRAT problem size can range from 5,000 up to 15,000 cargo units to be delivered during any one analysis. The number of available ships can range up to 1,000 ships that would be under MSC control in any given scenario. Ports could include 150 ports of embarkation and 80 ports of debarkation. The time horizon for analysis is usually 180 days, which means that multiple voyages are possible for each ship.

Despite the problem size and what I have heard here, I would like to get away from the heuristic approach. I know I am probably going to be locked into it to one degree or another. My thoughts are, after listening to you and Jai Jaikumar, who had the trucking model with the depots and so on, that perhaps we can break our problem into 1000 small problems and maybe then take a mathematical algorithm or something and begin to solve each individual problem, and then somehow put it all back together to come out with some kind of solution.

Now, that is just a thought. I will tell you where we are. Right now we are in the process of identifying all the data, the real scope of the problem, putting it all together. See Figure 10.6.

The description of cargo delivery requirements includes:

- . cargo characteristics such as type and amount
- . port of embarkation
- . port of debarkation and
- . earliest availability and required arrival date.

The description of ports includes:

- . port characteristics such as throughput, draft, available equipment
- . list of ports of embarkation
- . lists of ports of debarkation

The ship data base includes:

- ship characteristics such as size, speed, and cargo capability
- ship location and availability.

Ship routing data includes such data as the available shipping lanes which can be traveled and whether ocean clearance has been given for specified lanes. In addition, special shipping strategy factors include convoying, ship or port attrition, cargo unit cohesion (i.e., certain cargo should be shipped together), and resupply requirements.

We are working on the functional descriptions right now in order to come up with a new system to replace our present system, and I would like to open the floor to discussion, if anybody has any thoughts.

MACLEAN: In our previous workshop session we were discussing somewhat this kind of problem, and although you say you want to get away from heuristic solutions, it does appear that that is where you are going to start. The biggest difficulty appears to be defining the problem, and until you can put together the whole problem from the initiation of it through to the solution, you really have no way of identifying the bits and pieces of the compatibility requirements for the segments that are necessary in the final description of your problem. And, although you may not want to use it as a means of solving your problem downstream, if you can carry the process through you can at least then start attacking the individual aspects of it and trying to get a model for them to see if you can generalize the problem that way.

As far as I am concerned, that is a good, sound engineering approach to the problem. I think, that from a mathematical point of view, that is probably what you are going to have to do anyway, and it would seem to me that you are probably going to end up with the same kind of problem that Jai Jaikumar did in saying this is the biggest problem we can handle. We can handle only 100,000 variations now.

Well, how many variations are you going to have? You may have many, many more. It is a matter of the ability to exercise these if you are going to try to optimize.

GALLAGHER: I agree in principle. There is one aspect to identifying the problem. The problem is much like the problem that we had with the containers in putting them down; it is a moving target. One of the problems we have which you see here in Figure 10.6, strategy factors (down at the very bottom), is a very elusive problem to identify because it changes on a yearly, even maybe on a scenario, basis. You may change the complete concept of doing operations depending on what part of the world you are going to, so that your problem will be quite different, and you still have the same number of ships.

You still have the same amount of cargo, and now you are going to have to somehow deploy it in a new manner. And, as you alluded to, it becomes a very complex thing to put your finger on--exactly what the problem is.

MACLEAN: I think you have to classify those problems, though, as Tom Baker pointed out yesterday. Depending on what the constraints of your problem turn out to be, the solution technique may have to be sizably different. And, if you are going to change your strategies, you are putting different constraints into the problem that may preclude the applicability of a particular solution technique.

GALLAGHER: That is a key point. That is exactly what happened to us with SEACOP. As you recall, it was not an interactive system at all, it would batch feed at the front and give you a solution at the bottom, and as strategy changed and our objective changed, we could not alter it. You would have to front-load it and get the results at the end. And they were no good. They were not applicable to the situation that you were handed, and I think that is a good point, that somehow when we get done with this it is going to have to be partitioned, so that we get enough interaction to enable us to respond to a very elusive set of objectives.

WEBSTER: One of the things you have to keep track of, following up what Walt Maclean said, is that you have to be able to define your problem properly. Then you have this collection of tools, some of them heuristic, some of them exact solution methods of one sort or another. I do not

think that there is anybody in this room who does not know at least enough of both parts to be able to look at a given problem and make a judgment as to which one is most likely to work.

I think it is likely that both approaches have drawbacks and advantages, and if I may be so bold, I think we may be making a mistake to lock in on the target of having a mathematical solution, because it is going to cost a lot. It is going to cost you in your problem definition. It is probably going to cost you in how much complexity you are able to handle.

If there is one thing that heuristics can do, it is handle increasingly complex problems at minimal increase in cost. If you have already taken it on the chin to go heuristically, then additional complexity can be added without great difficulty. Mathematically exact solutions are often very sensitive to that, so that is a cost that you have to watch out for.

My suggestion really is that you remain flexible and not narrow in on either one or the other. They both have their disadvantages.

YOUNG: I guess one problem with the heuristic techniques, and this is the experience that Larry has already had with SEACOP, is that the solution just is not very good sometimes, and hand scheduling can do better. And in this case, he is actually manipulating the data in order to get the results.

Now, there are a number of heuristic approaches you can use. One is just some kind of logical scheme like that which Carroll Keyfauver has. You could have weighting factors with which the user can interact, or ranking schemes with different weighting factors, or some kind of empirical analysis of which weighting factors are best. But there is this real issue with heuristics of just how good the results are--you are always fixing up the problems that come out of your heuristics, and maybe you want to address that.

KEYFAUVER: You said something that I think is important. You are saying you are able to do it by hand better. I think you need to figure out how you are doing it by hand and make that your heuristic. Really, in a lot of ways that is a way to build a heuristic; find out what the mechanism is that you are going through to get a schedule that you are happy with and then automate it.

YOUNG: One of the aspects of that is, perhaps, providing explicit user interaction to allow them to play with the heuristics. I did not notice that in your model. I wonder if that might be a valuable tool to use also.

KEYFAUVER: There are user controllable parameters in heuristics that allow the user to play with them and get some idea of the sensitivity and whether or not they can be improved by playing with some of the parameters.

GALLAGHER: I would like to make a comment on that while we are on the subject of interaction. The interaction is great if you are smart, and if you are the person who has been handling the data base, and you are the person who has been working with it for a long time, and you are pretty sure of the results you want out of it; that is tremendous. Sure, I can go in and they can call on me, and if they want results, I can do it. I can load the SL-7s for them, just like that. I can do it, do a multitude of things with it, but it is not worth anything for Joe Blow who sits down and begins to run a war plan to come out with the first cut.

That is what is wrong with our system. It takes a tremendous amount of interaction, which I just have been recently developing the expertise to do by handling the data base, and so on. But if you get it too interactive, I do not have anything. I am back to a person sitting down there, and if I put eight people down there we get eight different results. And I do not have even the slightest idea which one is best.

Now, when we say we do it better with hand scheduling, that is not really the truth. If it was not for the computer that went ahead and scheduled 70 percent of it so we would have to deal with just the other 30 percent, if it was not for the computer that told us where our ships are going to be, when they were going to be there, we could never do it by hand. But once the results come out, it takes a lot of hand massaging to get what we would call an acceptable product. And to us, an acceptable product is a full ship sailing to a port.

MENTZ: You said something just now that I need an explanation on. You said that if eight people sat at an interactive session and came up with eight different solutions to a problem, you would not know which one was best. Can I ask you, wouldn't there be an objective function? Wouldn't there be a quantitative assessment of how well you met the objective function within some resource allocation, and wouldn't you know which one of the eight was best?

And then, furthermore, if all of the eight were very close, wouldn't you have to guess that maybe you were getting close to a "best" solution? If the eight were all far apart and one was by far the best, and the next one was far away and they were all spread apart, then you might guess that this is a wide-ranging problem set that perhaps you should do some more work on before you are comfortable with a solution.

I did not understand your comment that if you got eight different solutions, that would be an uncomfortable situation. I think that is a good situation and can lead you to an understanding of how well you are doing.

GALLAGHER: Well, if in fact my model is good enough that I could come up with eight different solutions, and one of them would be better than the others, then I would say the answer to the question is yes. The problem I have with the way the current model operates is that if I come up with eight solutions, all eight of them are likely to be wrong.

MENTZ: What does wrong mean?

GALLAGHER: I did not lift all the cargo and meet the RDDs.

MENTZ: It did not meet the requirements?

GALLAGHER: Right, it did not meet the requirements.

MENTZ: How do you know there is a feasible solution?

GALLAGHER: Because I can sit down with the analysts and go back and load the ships up properly, get the cargo that was missed, make the judgment call on the fact that I will overload my ship ten percent in this case and I will put a container on the deck of a break bulker in another case, and so on. I can go back and load them all and get my ships there. And after about

a three-week effort, I can then sit down and type up my new schedule. That is where it exists today.

So, yes, hand scheduling does it better. But it does it better because it supplements the job of taking those approximately 400 ships and loading them properly and sailing them properly.

MACLEAN: In my discussions with a fellow by the name of Seaver at EXXON, this is essentially the problem they deal with. The vast majority of their work can be done rather quickly and effectively, although I do not know whether the vast majority means 80 or 85 or 90 percent of it, through a quick, concise, manual operation and this can be routinized fairly effectively.

The real difficulty is the last five, ten or 15 percent of the work. And whether Tom Baker agrees or not, Seaver, at least, is at a point of despair trying to sufficiently well categorize and analyze the varying nature of that last small part so as to be able to handle it with an automatic routine.

Now, I have a feeling that is very analogous to the situation you are pointing to, and I think what Jai Jaikumar said the other day was that it appears that in these complex problems the objective function is relatively flat. So, where he feels he is reaching optimality within one or two percent, if you can get yourself a handful of solutions that are fairly close to it, you may have a good measure already as to where your optimality is and whether, in fact, you are satisfactorily close to it.

Another point I would like to make is that nothing is absolute. This is one of the things that I do not understand about the RDD and its implications, because the measure of effectiveness of an operation is not that you met everything absolutely, because nobody needs anything absolutely. It may be that therein lies the flaw in the analysis of the problem. One has to be able to determine what the effectiveness is of almost meeting something such as the RDD.

GALLAGHER: It is done this way, however, and it is a real judgment call, and it is not done by us. That final evolution is done after we get

through. We rarely make all of the RDDs, even after I scream and yell. We get everything picked up in the schedule, but there are usually some RDDs we cannot make. Ninety-nine times out of 100 there are RDDs that we do not make and we go to what I will call a manual effort. One goes to a Phase two conference off in Florida and the world descends and they make a value judgment on those RDDs being late; whether or not they can live with a particular gap. If they cannot live with it, then a priority is assigned to it and we will drop something else out. They will pre-position some force or they will do something different, so that with the assets given (remember, I am working from constraints on assets) we can then make the readjusted RDDs.

MACLEAN: But that suggests that this priority structure should be in better definition earlier in the game.

GALLAGHER: Or I should have the ability to put it into its proper place.

KEYFAUVER: Part of the problem is that the RDDs are not independent of the transportation system, not really, not ever. If they were independent, they would all be one day.

They start out with a guess as to when they can get the unit there, and it is a guess they think they can live with. So, in some sense, there should be an iterative process going on to improve the RDDs.

As CDR Gallagher has said, in some sense there is. When they cannot meet an RDD, they call a conference.

GALLAGHER: This is the position I want to be in with my modeling effort. With whatever system I end up with, what I want to be able to do is walk into that meeting, or face the world or the JCS, and know that I came as close as possible with a realistic optimal usage of the assets that were made available to lift the cargo.

Now, whether or not I made the basic RDDs, they will have to make the adjustments to these, but we cannot be making adjustments and hence making policy decisions. That is the hang-up I have. And the more deterministic and analytical I can get this, the better I am going to be

able to present the case that I need more ships, more money, more funds, pre-positioning of material for strategic mobility, and so on.

GARFINKEL: There are two quick points I would like to make. Number one, there is the contradiction that you are saying you would like to be as close as possible when the objective has not been defined. Clearly, it would be nice to have an objective, and then, of course, you need some sort of weights on infeasibilities. And then comes question number two of whether you want to talk about heuristics or optimization.

The question is, how do you define what Jai Jaikumar was talking about? Was that heuristic or an optimization? It was a heuristic that was based on optimization, and it seems clear to me that what you would like to have is something similar so you would be able to say "I am within this much of optimality."

To have that, first of all, you need to be able to define "optimality" when you are talking about feasibility. So you have to have either the weights or something else.

I would guess that you want to develop a system, at least a heuristic system that is based on optimality, and have at least a first cut at the objective function, and then be able to interact in the sense that when something comes out as the solution, then you say to yourself, "Well, this has come up or that has come up." and make a change to your data and go again. I think it has to be an iterative system.

GALLAGHER: I am in full agreement with you. If there is one statement I can stick to, it is that there will always be a requirement for me to load a ship manually. In other words, there is not any model that is going to put all of the right cargo in the right ship, or make a decision for me that I am going to overload a ship ten percent, or take that last piece of cargo and throw it on there and ship it.

When you get through with your model, you are going to have to go in there at some point in time, when it has done its best job and tidy up.

GARFINKEL: You indicated 180 days earlier. What does 180 days stand for?

CALLAGHER: That is the time span of the exercise, and it means a ship can make multiple trips. The faster the ship, the more trips it can make.

GARFINKEL: How often does this model have to be run?

GALLAGHER: Currently, we have 12 war plans with 12 different scenarios. That could evolve to 30 scenarios, depending on how or when we decide where we want to fight a war and under what terms we want to fight it. At one time the model SEACOP was built to be run once a year for one war plan, and that was the reason it was batch loaded. Yesterday I ran four war plans in a day with the modification that was made to SEACOP.

So how often do I have to run? At least three or four times a week as it stands now. As you get closer and closer to the solution and you begin to examine more options to get what you might call the optimum output, you run it one, two, three, four, five, six, seven, eight, nine times.

GARFINKEL: The reason for that question is that this tells you how much time you have to run the problem, and that, fundamentally, tells you how large a problem you can run. Jai Jaikumar could only treat 100,000 routes. Perhaps you can treat a million routes.

GALLAGHER: My present SEACOP will select ten ports. It goes through computations. At one time it was taking 48 hours to run one war plan. Now of course we have done a lot of things to it. We are down to about four hours, and that is still a long time when you talk about computer time.

WEBSTER: I wonder if you could comment on the following: no matter what system you use you are going to have some kind of post-processing massaging, and the final solution as a result of that massaging is now of questionable optimality, I think, even if the solution you have to begin with is optimal.

Now, it seems to me you are giving up on optimality already by this process. Your choices of overloading a ship or putting containers or whatever on bulk carriers, if it is as you describe, rather arbitrary, then you have departed from your goal.

BRIERRE: You will never have a model that contains everything.

YOUNG: I think it is a good point that there should be some kind of cost associated with failure to meet delivery dates. There should not be hard and fixed constraints so most of the time the problem is infeasible. The reason they do so much post-processing is that the heuristics are not adequate.

WEBSTER: It sounds to me like you have the wrong constraints. If you allow ships to be ten percent overloaded, and you do not crank that in as a constraint, then you are unfairly judging your problem, and the solution method. Maybe that is a concern of the data base that CDR Gallagher has referred to. If you allow the ships to be ten percent overloaded, why not inform the computer program of this?

GALLAGHER: I could list for you 47 similar points that I could try to put into the heuristics of SEACOP.

YOUNG: Such as "Do you go to multiple drop points? Why do you sail ten percent? Why do you do this, why do you do that?"

GALLAGHER: Part of the problem is that the logic of the system is not very well documented. It is really not worth going through tremendous reprogramming efforts just for the schedule part of SEACOP. That is why we are here, and we are trying to define the problem and go to a new system.

Instead of going back to do all of the iterations to answer the questions I ask myself, what I have been successful in doing is solving some of those problems by messing around with the data base. If I find something, then I can pull it out and hand-schedule it.

MAYS: I think I heard yesterday that there is high priority on some of these things. When do you have a requirement to have this new system up and working?

GALLAGHER: I expect to have this up and working in 1983.

MAYS: Good luck.

GALLAGHER: It does not have to be totally done. If I can define the overall problem and start attacking the smaller problems one at a time and come up with some modeling to take each problem as I come to it, I might have 30 of the 500 problems solved by the end of the year. And it would seem to me it would be a continuing process, to continue further defining the system.

KEYFAUVER: At what level of detail are you loading ships? Do you have some gross numbers to use or are you loading individual pieces of equipment on each ship to estimate its payload?

GALLAGHER: I am loading each ship to its individual characteristics of cargo. I am using the characteristics that define its lift, its wheel, and so on.

MACLEAN; Are you worried about such operational matters as stability and trim?

GALLAGHER: No.

MACLEAN: You do not know whether you have generated a horrendous stevedoring problem or a problem that this ship will have difficulty overcoming?

GALLAGHER: Yes, sir. I do to some degree. There are only certain kinds of equipment I will put on certain kinds of ships. Each one of the ships has a certain amount of cargo it will handle. Each one of my ports has a through-put limit based on stevedoring and berths and that sort of thing. That is all in the model.

MACLEAN: Do you differentiate between deadweight and volume?

GALLAGHER: Yes, we do. We have broken storage factors that are in there and a bunch of other stuff in there also.

One problem is that you never get a good shipping exercise. I can sit here and plan until I fall over, but we never get a chance to execute a plan.

MACLEAN: Don't you have an opportunity to exercise it on some of the MSC ships that are actually running.

GALLAGHER: That is a good point. No.

MACLEAN: Even in a small subset type of case? One would think it would be an excellent tool for trying it.

GALLAGHER: That is where the numbers come from. I load the ships with that cargo. That makes the cargo that goes on an MSC ship in an exercise of the model the best kind of cargo to put on that ship.

MACLEAN: Why can't you use some of the requirements for those MSC services as a means of testing in a smaller fashion whether, in fact, your idea will work in reality.

GALLAGHER: I guess we probably have one Sealift exercise a year, so we do that to a degree. These are usually one or two ships loaded.

JESMER: There have been studies made of putting various type units on ships, engineering studies, and in fact, the ships were actually loaded with military equipment. We do have a feel for cargo densities and volumes.

MACLEAN: But he needs to have an exercise of this model to find out from scratch. You have got the facts in place and you have a tool that is going to provide the decisions for you. And you carry the thing through and then find out whether it works.

GALLAGHER: The ideal thing would be to take a canned scenario that I know the exact results for, run it, and then go ahead and run my planning data and see if I get the same results. I do not have that. I do not own it. It is just not available.

I guess you would have to get those results by people sitting down and hand-scheduling, and then you would still be guessing whether or not you would be able to have the ships loaded correctly. It would still be based on a person sitting down there. It certainly would be a good test.

MACLEAN: So you have a good validation problem that has to be faced up to, I think.

GALLAGHER: SEASTRAT modeling requirements are shown in Figures 10.7 and 10.8. SEASTRAT should schedule realistic cargo movements consistent with available ship characteristics, ship locations, cargo loading/unloading times, port throughput capabilities, and shipping routes. SEASTRAT should produce both summary feasibility results as well as detailed, ship-by-ship schedules. SEASTRAT should incorporate reasonable ease of use with practicality in input data. SEASTRAT should provide reasonable computation time, especially for high-level feasibility analyses. SEASTRAT should provide MSC user interaction for adding constraints and performing sensitivity analyses or "what if" evaluation. SEASTRAT should provide capability to address a large number of ships, ports, and cargos as required by operations plans (OPLANS). SEASTRAT should address special crisis features such as:

- . attrition
- . convoying
- . sanitized sea lanes
- . cargo cohesion
- . resupply.

SEASTRAT should accommodate statistical analyses of ship locations and other key uncertainties. SEASTRAT should determine degree of shipping feasibility/infeasibility (how late?) and identify areas of improvement. SEASTRAT should interface with land transportation facilities.

I would like to have Steve Young, who has been working with us, give you conceptual ideas for attacking this problem, breaking it down and partitioning it.

YOUNG: To start you modeling experts thinking, Figure 10.9 shows several generic methodology issues for SEASTRAT development. The first issue is level of detail, which can range from gross feasibility to a detailed, executable schedule. Gross feasibility answers such questions as: are there enough ships to lift the cargo in the time required (irrespective of details such as port throughput, convoying, etc.)? An executable schedule, on the other hand, is so detailed that the cargo and ship movements could be implemented in the "real world."

SEASTRAT MODELING REQUIREMENTS

- SCHEDULE REALISTIC CARGO MOVEMENTS CONSISTENT WITH AVAILABLE SHIP CHARACTERISTICS, SHIP LOCATIONS, CARGO LOADING/UNLOADING TIMES, PORT THROUGHPUT CAPABILITIES, SHIPPING ROUTES
- PRODUCE BOTH SUMMARY FEASIBILITY RESULTS AS WELL AS DETAILED, SHIP-BY-SHIP SCHEDULES
- INCORPORATE REASONABLE EASE OF USE WITH PRACTICALITY IN INPUT DATA
- PROVIDE REASONABLE COMPUTATION TIME, ESPECIALLY FOR HIGH-LEVEL FEASIBILITY ANALYSES
- PROVIDE USER INTERACTION FOR ADDING CONSTRAINTS AND PERFORMING SENSITIVITY ANALYSIS AND "WHAT IF" EVALUATIONS

FIGURE 10.7

SEA STRAT MODELING REQUIREMENTS (CONT.)

- PROVIDE CAPABILITY TO ADDRESS A LARGE NUMBER OF SHIPS, PORTS, AND CARGOS AS REQUIRED BY OPERATIONS PLANS
- ADDRESS SPECIAL CRISIS FEATURES SUCH AS
 - ATTRITION
 - CONVOYING
 - SANITIZED SEA LANES
 - CARGO COHESION
 - RESUPPLY
- ACCOMMODATE STATISTICAL ANALYSES OF SHIP LOCATIONS AND OTHER KEY UNCERTAINTIES
- DETERMINE DEGREE OF SHIPPING FEASIBILITY/INFEASIBILITY (HOW LATE?) AND IDENTIFY AREAS OF IMPROVEMENT
- INTERFACE WITH LAND TRANSPORTATION FACILITIES

FIGURE 10.8

SEA STRAT METHODOLOGY ISSUES

- LEVEL OF DETAIL - GROSS FEASIBILITY VERSUS EXECUTABLE SCHEDULE
- SCHEDULING APPROACH - HEURISTIC VERSUS OPTIMIZATION VERSUS HYBRID
- APPROACHES FOR IMPROVING TRACTABILITY
 - AGGREGATION
 - PARTITIONING OR DECOMPOSITION
 - SUCCESSIVE APPROXIMATION
- SPECIAL FEATURES
 - CONVOYING
 - SHIP LOCATION STATISTICS
 - LAND TRANSPORTATION INTERFACE
- THE ABOVE ISSUES WILL BE ADDRESSED IN THE WORKSHOP

FIGURE 10.9

A second issue is the actual scheduling algorithm to be used in SEASTRAT. Heuristic methods use reasonable logic and decision rules to schedule ships and cargos --the problem comes in defining what is "reasonable." An optimization approach uses mathematical techniques to determine the "best" solution; however, such techniques may not be computationally feasible for large problems. A hybrid approach could combine both heuristic and optimization methods.

An additional consideration is the need for modeling approaches which improve computational tractability. Aggregation can be used to combine small elements such as individual cargo, ships, or ports into a larger element, which can then be treated as a single unit without altering the basic features of the problem. Partitioning or decomposition techniques can be used to structure a high-level, large model into smaller submodels which can then be solved nearly independently. For example, individual ports could be partitioned into port regions. Approximation can be used to model complicated relationships in simpler terms, with subsequent improved approximations made as the solution is obtained.

Finally SEASTRAT has a number of special features which should be addressed in the model. These features include the need for convoying, the statistical nature of ship locations at the starting point of a mobilization requirement, and the need for interface with land transportation facilities.

One of the big issues that we are thinking about is this question of level of detail. That is a problem because MSC needs both quick turnaround, gross feasibility, and very detailed, executable-type schedules. And there is some need to have a model that is capable of the full range. That is a big difficulty.

And, of course, the other thing we have been talking about is this heuristic versus optimization versus some kind of mix, as Jai Jaikumar presented yesterday.

But one of the big questions is how we can improve the tractability of this monstrous problem. There are several standard approaches that people use. One is aggregation, combining cargoes. They may have 50 cargoes at a port that are starting at the same place and going to the same place. If we combine them as one big cargo and treat them as one, as an aggregate, it is much simpler.

Other things are partitioning or decomposition techniques, whereby you might combine several ports in a region and treat them as a subproblem. Or you might use some kinds of approximations.

In your experience, Carroll, have you used any of these kinds of techniques to try to simplify the problem a little bit, break it down?

KEYFAUVER: We deal with the aggregation problem. We are not working at the great level of detail you are, and generally we can reduce 10,000 to 15,000 units to probably 3,000 to 4,000 by aggregating those things that are going from the same place to the same place at the same time and that are of similar composition.

YOUNG: So aggregation has been a useful approach for you?

KEYFAUVER: Yes, and it is going to be one of the critical factors in the solution time constraint that you have here.

YOUNG: Have you tried any partitioning or decomposition in what you are doing? For example, you have a land problem and a sea problem, so the overall problem could be partitioned, perhaps.

KEYFAUVER: In effect, we do that, and I am solving both problems. I am solving each individually, and then using a mechanism to revise the decision process, as I discover that I would be better off switching modes for a particular unit.

YOUNG: I think we also had an example of decomposition in Jai Jaikumar's approach, using Lagrange multipliers to develop subproblems which he solved separately.

KEYFAUVER: Decomposition is an optimization technique that you should be considering for this kind of problem, if you can break it into problems that are small enough to solve individually in reasonable time.

One of the big problems is what kind of objective function you use. Probably more than one, in succession. You have a strong emphasis on feasibility. That would be one possibility.

There are other objectives that you should consider, even if you get a feasible solution or a solution which in some sense minimizes the infeasibility. There are other objectives that you have in the deployment in terms of fragmenting each individual unit as you deploy it. If you need one piece two weeks earlier than another piece, you have a real problem.

YOUNG: Bill, did you have any kind of optimization or partitioning techniques in the container loading problem?

WEBSTER: Well, we tried. We have tried a lot of different optimization techniques. The simulation is partitioned in a sense. I have developed a module which loads a container ship at a port, so it removes the containers that get off at this port, adds the containers that get on at that port and ranks what happened. This module is called over and over again as you go further downstream. So there is a kind of partitioning there, but it is perhaps a recursion rather than partitioning.

YOUNG: It is applied to each port. You separate the ports essentially?

WEBSTER: Each port fundamentally looks like each other port, and the important part of the container ship loading problem is whether you are doing something in the current port which will cause you grief three or four ports downstream. The way I deal with that is, I try lots of things at this port and then follow those through and constantly rank and weed out those that are absolutely bad. I guess it is a kind of partitioning. I would call it recursion.

KEYFAUVER: We do something I mentioned yesterday to which you should give some consideration. It does not really fit into any of these categories specifically, but it is in the way of partitioning. It is to limit in time the scope of the individual problem to solve. I see no real good reason why you have to solve a problem that extends over 180 days to get your immediate answer.

There are probably lots of advantages in solving the problem for more limited periods of time, intersecting periods of time, because something that is late early in the problem is a lot more critical and generates a lot more concern than something that is late on the 180th day, when you could not care less.

You have a number of alternatives that you can look at in that kind of analysis.

YOUNG: That is a decomposition in time.

KEYFAUVER: There is a real problem with trying to measure the value of units over time.

WEBSTER: That is, incidentally, an advantage of the heuristic approach. With the strict optimality approach you are going to get the answer for the whole 180 days, and you have to take it on the nose for computing the whole 180-day answer if you get anything; whereas, heuristically, as soon as you hit this flag that you were talking about, "Gee, I missed this RDD on day 25," then you can stop, juggle, and do something without paying for all of that extra computation.

FRIESZ: My question is mentioned at the bottom of Figure 10.9. I would like to ask how the interface between the land loads and the ship scheduling and routing problem occurs.

It seems to me in a crisis many of the commodities that you have to move are not in ports. They have to be moved from places interior to the country, and that is a whole other problem, an additional level of complexity.

GALLAGHER: That is not really our problem, but it may well become our problem under the new merger. Right now their model takes it to the ports and dumps it there; I pick it up at the port. They use backward planning to get it to the ports on time. They start with an RDD date and back it up, take it all the way back to, say, Little Rock, Arkansas, put it on a truck and get it to the port, all based on the assumption that I will have a ship there that will deliver by the RDD.

FRIESZ: Do their estimates for the amounts of these commodities in the ports include any sort of confidence intervals? It seems that it would be hard to make that exact determination.

GALLAGHER: I do not know.

KEYFAUVER: One of the points you are getting at involves the ship location statistics. I want to point out that these ships are moving around all over the world.

You have two alternatives: you can either take a snapshot of where they are at some point in time, and use that as an input. Or you can use some kind of statistical approximations with some confidence intervals. Within that, you have some real problems with capability. If you get a feasible solution, then you have a problem of how to place confidence in that solution, given the statistical location of the ships that you use as input. If they are not in those locations, but are in some other kind of configuration, then you may not really have a feasible solution.

VOICE: The complete problem is the land problem and the ship scheduling problem simultaneously.

JESMER: The complete problem is to solve the land problems, ship problems, and air problems. We have airlift working, and there are days when ships are not available to move the cargo but airlift would be available. So you have slack both ways. The airlift is not utilized 100 percent of the time, or even 100 percent on any given day.

VOICE: So even if we have optimality models for land, air, and sea, the way they are being applied is heuristic, anyway. We are using them sequentially.

KEYFAUVER: You are truly suboptimizing.

VOCE: I think it begs the question of how much energy you put into getting an optimal solution.

YOUNG: From Larry Gallagher's point of view, he has the cargo at these ports, and he has got to come up with some kind of algorithm to schedule it. He can tell the land people that he has problems here and there, but he has to be able to justify telling them, and justify saying he just can not meet the schedule. It is true that it is suboptimum from the total point of view, but he has to be able to justify that statement in terms of his portion of it.

GALLAGHER: I would put forth the following thought: if we try to have a system that will take it from the very start, and take it to the end, we will never get off the ground. If we are ever going to solve the problem it is going to be solved in pieces, and maybe at some future time some big distributive type model and solution might be there.

But until such time, it is going to be a piece-by-piece solution. We are going to try to optimize, though, at the level where we can meet the mission, and to optimize at the higher level is a long way off.

GARFINKEL: It seems to me that we get back to the question of optimalities, since these delivery dates are things that people perhaps can live with. It could be the case that a suboptimal solution might be better than an optimal solution in the sense that somebody could also live with something else --but something is more important than something else, and if you can get the important thing there a little earlier at the expense of the less important thing a little later, you may not, in fact, meet the delivery dates given to you, but that solution might, in fact, be better in terms of the overall value of the solution.

Again, it seems it is critical, somehow, to get values associated with those delivery dates in order to say which is better. Optimality is not even defined, as we now understand it.

GALLAGHER: Absolutely true. An early RDD is much more important in most cases than a later RDD because later on you get down to 180 days, and you go into resupply (unless it happens to be bullets). But there are time values, there are unit values, there are material values, that could be assigned, as we were discussing, to begin to work the problem better than we are doing now.

KEYFAUVER: Getting two people to agree on those values may be difficult.

YOUNG: I think that is where user interaction comes in.

WEBSTER: Do you have any feel for how many solutions you have that are close enough to being optimal?

I noticed in the container ship problem that there is essentially an infinite number of solutions that are so close together, in terms of optimality, that selecting any one is okay. The actual selection is done on bases which you cannot crank into a computer program.

Do you have any idea whether, in your problem, there is also a whole collection of solutions that for all practical purposes are equally optimal?

GALLAGHER: Yes, there probably is, particularly if you are speaking of an individual ship, for utilization, or if you are looking at the overall plan. The way we select ships and load them and put certain equipment on them, I am sure that there must be an infinite number of other choices that could be made, and which would yield equally good solutions.

WEBSTER: Then the obvious statement is, isn't it sufficient to find one of the bunch, rather than to find the best of the bunch?

GALLAGHER: Very definitely.

KEYFAUVER: I might add as an example of that something we did eight or ten years ago, when we were using linear programming to do this kind of

optimization on a somewhat more aggregated problem. We would get an optimum solution using the objective function for some supposed measure of maximizing capability, and then, having obtained the optimum linear program solution, try to get a binary integer programming solution, to get a solution which assembled convoys. Each convoy had to be scheduled, and running the thing through for hours we ended up with a solution, the first feasible solution, which was within about one-hundredth of one percent of the original objective function value, consistently, and at great expense.

That is an indication that there is a broad range of solutions that is probably quite similar to what you would consider optimal.

GALLAGHER: Yes.

YOUNG: Certainly that tends to be a feature of an integer programming problem.

The broader question is really: is there hope for a mathematical programming technique, or are we really just going to be resorting to heuristics in trying to solve these types of problems? (Also see Figures 10.10 and 10.11).

Is there any concensus?

GARFINKEL: When you say "mathematical programming," and then you say "heuristic," it is not necessarily contradictory. I would guess there would be hope for something like that on the order of a hybrid technique.

SEA STRAT SCHEDULING APPROACHES

- HEURISTIC TECHNIQUES
 - LOGICAL DECISION RULES
 - WEIGHTING OR RANKING FUNCTIONS
 - SEMI-EMPIRICAL
 - OTHER
- OPTIMIZATION TECHNIQUES
 - LINEAR PROGRAMMING
 - INTEGER PROGRAMMING
 - DYNAMIC PROGRAMMING
 - OTHER MATHEMATICAL PROGRAMMING
- HYBRID METHODS

FIGURE 10.10

SEASTRAT APPROACH

SUCCESSIVE PARTITIONING

- PARTITIONING OF CARGO/PORT/SHIP DATA BASE
WITH USER INTERACTION
- DECOMPOSITION OF HIGH-LEVEL, GROSS FEASIBILITY
PROBLEM INTO SMALLER, MORE DETAILED AND TRACTABLE
SCHEDULING PROBLEMS

FIGURE 10.11

Discussion Group 5 - Exact

Solution Approaches

Leader: Nicos Christofides

Department of Management Science
Imperial College
London, England

CHRISTOFIDES: We are supposed to discuss the possibility of solving shipping problems, i.e., routing and scheduling problems, in an exact way.

You may have gathered from the discussions over the last couple of days that there are a number of exact approaches to these problems. But before going on to discuss some of the approaches, I would like to just point out a few things about what we actually mean by "exact." And I would like your comments on what we can reasonably assume to be allowable and what is not (see Figure 12.1).

For example, what about aggregation, either in terms of time (that is, taking time slots which are not infinitely divisible, but using half-days or hours or minutes) or perhaps in terms of distance, by putting together ports of call that are close to each other? Or, if in some exact models the types of ships rather than the number of ships is a more important factor, how close do two ships have to be to be called the same type? There are aggregation problems of all kinds.

There is a problem that comes up from yesterday's presentation by Jai Jaikumar. If the routes are to be partially enumerated and then a subset of them chosen to be the solution, then what detail does one go to in partially enumerating routes? And what guarantees can one give in saying that a problem solved optimally with certain routes considered is within some error of the problem solved optimally with all the routes enumerated?

This is again a problem in the preformulation state that has to be addressed before one is looking at solving a model exactly.

PROPERTIES AFFECTING THE "EXACTNESS"
OF A MODEL

1. AGGREGATION
2. PARTIAL ENUMERATION OF ROUTES
(FEASIBILITY PROBLEMS)
3. INTERACTION OF SUBPROBLEMS
4. ROLLING HORIZON - END STATE UNSPECIFIED
5. UNCERTAINTY

FIGURE 12.1

It is clear that one cannot look at a ship routing or scheduling problem in its totality because, from what we heard yesterday, they seem to be extraordinarily involved and complex problems. So one is naturally interested in breaking them down into subproblems. And the interaction of the subproblems will obviously affect the global solution to the entire problem.

So one has to decide how big the subproblems should be. Jai brought up the same question yesterday. The problem exists generally in all kinds of modelling and has nothing to do with shipping. Obviously, the bigger the subproblem, the fewer interactions you're going to have between the subproblems and the less you'll have to worry about. But then you have to worry about solving large subproblems. So the disaggregation question of breaking up a complex system into a number of simpler ones produces these problems as far as exact solutions are concerned.

There is the additional problem of the rolling horizon. Personally, I was only involved in one ship routing problem. I have extensive experience in road vehicle routing and scheduling, but only with one problem in ship routing. It had a characteristic that gave me more trouble than anything else, and it is one that does not appear in road transport problems. It was the property that there were no definite ends. It was a monthly problem of distribution of chemical products from refineries to depots. The horizon was one month in this case, but there was no specified end state so I could not say that at the end of a month these ships were going to go back where they were at the beginning of the month.

That characteristic produces a lot of difficulty in exact methods for solving these problems. In fact, it produces a lot of difficulty in just formulating these problems.

Obviously, if you have a long enough horizon, the end-state is not going to matter very much. Whatever the end-state is at the end of a

one-year rolling horizon, it's not going to affect the solution for the first few days very much. But what is the length of this horizon in order to have some confidence in the results that the model produces in the short term?

All these problems have nothing to do with uncertainty. I would like to emphasize this. They are problems that exist in a deterministic sense. In addition to these, you have a whole host of problems involving uncertainty, and rather than list many of them, I just used the single title as point 5 in Figure 12.1

Perhaps we can start by having a discussion. I would like your comments on what you think about producing some acceptable level of exactness in a model.

BALLOU: Could you make your comments again on partial enumeration of routes?

CHRISTOFIDES: Looking at any routing problem is a two-stage process, where in the first stage you enumerate a set of feasible single routes, and in the second stage you choose a subset of these routes. One asks first "What single route could be operated feasibly?"

That, in itself, can be a difficulty in some situations. There are even cases with a single vehicle and a given set of customers with time windows during which the visits have to be made, and so on, where it is not simple at all to find a feasible solution, let alone an optimal solution. So there are problems involved in the partial enumeration which become very serious if there are difficulties in finding a feasible schedule to meet the requirements.

Suppose we ignore feasibility problems. Suppose it is easy to find a feasible solution to this problem; then we look for an optimal solution. We generate a lot of routes, say, 100,000, as mentioned yesterday. And suppose, despite the fact that the problem is very large, we can solve it optimally.

The answer that one gets is not optimal to the problem as originally stated. Simply because you enumerated 100,000 routes instead of the billions; 100,000 is a very small number in comparison with the real number of possible routes that exist in this type of problem.

So the question arises as to how far one is from the solution to the original problem before you decided to treat it as a two-stage problem; first enumerating a set of routes and then choosing the best combination of those routes. It is a question of having some guarantee that at the modelling stage you have not introduced a significant error.

JAINUMAR: In terms of aggregation, I see two types. One is where you are aggregating to reduce the size of the problem. The other one concerns the data; if you aggregate, they are more certain in some sense. The accuracy of the data improves with aggregation. If it is the former kind of aggregation, one can try to analyze it. But if it is the latter, I do not know the techniques to use. When do you aggregate if the data are not accurate?

CHRISTOFIDES: It is not for accuracy. In road vehicle routing, if you have two customers in the same street and they are a few hundred meters apart, you may well consider them as one customer, depending on the size of the problem. If the problem involves a daily call to 2000 customers, for whom the number of distinct locations is only 200, clearly you would prefer to deal with the 200 rather than the 2000. So that is aggregation to reduce the size of the problem. Of course, it also reduces the input data requirement, which is perhaps equally important.

JAINUMAR: If a route takes 8 hours, that is 8 hours which you think it is going to take. But because of uncertainty, it could be 7-1/2 hours or 8-1/2 hours. Now should you model in terms of hours or in half-hour blocks? How do you define resolution in these problems?

CHRISTOFIDES: I tried to take uncertainty out of our discussion. Very often there is a confusion that the problems of indecision come because of uncertainty, and I do not believe it is so. I think there are lots of problems of indecision that have nothing to do with uncertainty, that uncertainty is an additional problem which is blamed for a lot of things.

MACLEAN: It may be that the information that is available for decision is inadequate or incomplete. Or perhaps it is inaccurate, but it does not mean that something is uncertain in terms of generating it.

CHRISTOFIDES: That is right. It is simply not known. For example, what time the vehicle has left is something that can be known with certainty.

RONEN: Referring to the last point, uncertainty, which is a major obstacle in ship scheduling, one way to deal with it, I think, is official: to reduce uncertainty by shortening the planning horizon. If you look, for example, at what is going to happen only during the next week, you more-or-less already know what are the cargoes that are available for loading during that period and what are the ships available to carry the cargoes.

So in using a short planning horizon, you have one way to tackle uncertainty. But then you have the problem of how you relate one plan to another, the rolling horizon problem.

CHRISTOFIDES: Certainly, operational data are more accurate in the short-term than in the long-term. I think it was mentioned yesterday that a lot of the optimization savings in these problems have to do with a good coupling of one period with the next when considering the problem as a multi-period problem. This is certainly true in road transport. In fact, one of the most important problems in road transport is the allocation of calls to dates. When do you go to different places, as opposed to which sequence do you follow?

I would suspect it is the same in ship routing and scheduling. The time schedule, the interaction between one period and the next, is an important source of savings that is to be exploited.

ENGLISH: Do I understand that this dialogue suggests that there is some hope for an interactive scheme where you can essentially redefine the conditions periodically by operator input and rerun the problem immediately after, that kind of thing?

RONEN: There definitely is hope for this.

ENGLISH: But how about from a practical point of view? That is one of the things that is hovering in the air about this discussion, if I am guessing right.

RONEN: The problem is how to tie the single period into the multiple-period problem.

CHRISTOFIDES: There is certainly a problem of having to decide what happens at the end of the period that you are considering. If you are considering a single day, clearly it would be ridiculous not to take into account where the ships are going to be at the end of the day. Where does this position stop? Would it be reasonable at the end of the week; to ignore the end-state at the end of the week or at the end of the month? It is an open question.

RONEN: The problem is that in shipping your decisions are effective for a much longer time than in road transport scheduling. The decisions you make in a week probably will affect you during the next month, at least.

GREER: The big asset is the ship. So its position, its use, matters, because if you say you can dispose of the ship in two years, then you do not really need to look ahead more than two years. But if you have to have the ship for five years, you had better look at it over five years. All of a sudden, after two years or after three months, you may solve the problem and you are in trouble because you do not have use for the ship. That is what happens in shipping companies, of course. They have too many assets all of a sudden and they do not have the business.

So you have got to look at the long-range problem as well as the short-range solution.

ENGLISH: That might be a forecasting problem that transcends the routing and scheduling. And so if it is one of the go-to-war type problems, if I can inject another point of view into this, it is the more deliberate look at the thing.

But it seems to me that, once again, it is trying to understand the guys who have actually been in there, into the thing with their sleeves rolled up, to ask what they are telling each other. With regard to deliberate planning, where you really are planning in the abstract to try to establish planning values, planning decisions, is it possible that you could have a module with regard to the computer models that I think are, once again, implicitly a feature of this discussion, with a model off to the side that introduces the uncertainties, that generates these, and essentially, the more exact solutions would run to a certain point in time and then the uncertainties would be injected into the modelling on a random number basis and you would stop and you would get a new set of conditions and then you would run again with the change in the conditions?

Obviously, you would take the end conditions at the prior time interval as a point of departure, but these would be perturbed. For example, ships would be sunk, ships would be delayed, any of these things. But this would be injected as a perturbation of the final state of the prior time period.

Is this a little bit along the line of what the professionals in this field think about a way to inject uncertainty? And then obviously, if you go to the crisis situation, you would ask the computer to give you useful solutions over intervals of time. The commander who has charge of these assets in winning the struggle will ask, "what do I do next?" And he will ask the computer a lot of things. Then, it seems to me, in practice you operate similarly, only you do not draw these out of a module of the computation--you draw them out of the real world. You get reports and at certain intervals you interactively insert reports back in to the computer.

CHRISTOFIDES: Clearly, this is a plan which is not going to be applied in a vacuum.

ENGLUND: But is there room in the more exact approaches that you chaps are familiar with for this kind of interactive thing to provide both in the deliberate modelling and in the real world for injection of the results of uncertainty?

CHRISTOFIDES: Yes, I am sure there is room for that.

KASKIN: With respect to your point number two on partial enumeration of routes, are you implicitly assuming that the best method of handling scheduling problems today is by enumeration of routes?

CHRISTOFIDES: No. I was simply putting down a number of points that came immediately to mind from yesterday's discussions. Every one of those came up in some form or another yesterday.

KASKIN: Does doing that simplify the problem in that the routes do not have to be selected as part of the model?

CHRISTOFIDES: Yes.

KASKIN: Is that because there is no methodology today to optimally select routes?

CHRISTOFIDES: No, there is. There is certainly methodology which avoids problems of how to select routes. You do not have to consider partially enumerated routes. The formation of the routes comes partly from the formulation of the problem.

KASKIN: Is it known that if you add that extra dimension it often makes problems computationally intractable?

CHRISTOFIDES: In some problems it would, and in other problems it would not. It depends on what decision variable you use. A branch-and-bound scheme was described yesterday. The branches in the scheme have to fix some kind of variable, perhaps a new port of call in an emerging route. There are many different branching schemes for doing that kind of branch and bound. Some of them are computationally intractable and some of them are not. Some of them are perfectly feasible branch-and-bound schemes that avoid this problem of selecting routes.

KASKIN: Is there enough information about our problems for people to give a first impression about whether it would be or would not be a problem?

CHRISTOFIDES: I do not have any information myself. Maybe Jai Jaikumar would have some.

JAIKUMAR: We experimented with partial enumeration. You can add routes as you go along, but you keep them ready. But if the fixed cost is high, then it is worthwhile just generating the route when you need it.

So it is a trade-off, but it is almost a function, really, of the data. And in some instances it works out that you do not have to trade very much. Also, one of the things which we felt important in this particular application was that when you generate a large set of routes it gives the customer the feeling the optimal solution has been considered. In one of the applications, at one particular terminal which we looked at, the set of routes we had to look at was much smaller. And that is, in fact, optimal. You get the optimal solution. But the client feels that looking at only such a small set of routes may not be giving the optimal solution.

So there is a comfort index in looking at a large set of routes. A partial enumeration sometimes is not a subset, but a superset. You are doing more than you really need to do.

CHRISTOFIDES: There is an implicit meaning in enumerating a route. How do you enumerate a route? You start your route and you decide that this specific call is the first call in the route. (The other possible decision is that this is not the first call in the route, which is ignored and which could be considered later.) Then this other call is the second call of the route, or it is not the second call in the route, etc.

So in the sequence of decisions you have a route which you can either take apart and consider later, or you can consider now as an

element of the branch-and-bound. So it is difficult to consider these as two separate methods. It is a different way of searching a branch-and-bound tree by ignoring parts of the tree that you do not want to examine. But the relation of the route to the tree itself is more-or-less the same.

GARFINKEL: I think it is important to be clear on why one would be interested in enumerating routes rather than just using some other mathematical formulation. And I think the reason is you have a lot of feasibility constraints on the routes which may be difficult to formulate mathematically, which you may have to leave out otherwise.

Then it may be that the only practical way to have an optimization-based code is, in fact, to do that and then you can say once you have accepted a route that it is feasible. You do not have to worry about it in a mathematical statement dealing with the other aspects of the problem.

CHRISTOFIDES: I am not sure that is true because you are going to generate this route in some way. Nobody is giving you this route. You are going to have some subroutine that will generate it. This route is generated by some sequence of decisions so in the intermediate step you could easily say: "This is a partially formed route. Do I want to continue with generating routes that have this as a subset or not?" And you can perhaps eliminate a whole family of routes that you might otherwise have to generate, simply by considering the problem that remains having partially formed a route; rather than completing it first before you consider it.

GARFINKEL: I am not sure I understand what you mean.

CHRISTOFIDES: I thought that you perhaps implied that the feasibility of a partially formed route is harder to check than that of a completely generated route.

GARFINKEL: No. It may be that it is very easy to check, but it may be that you might not want to have this checking be stated in some

mathematical way. Instead you have checks which can be stated in some fairly vague way that the computer can check, and you can still possibly eliminate whole sets of routes. But you do not have to go through some sort of optimization routine to check feasibility.

SOLAND: I would like to say first that because of some of those points, particularly 1 and 4, of aggregation and rolling horizon, we might say that there is no such thing as an exact model, and therefore no such thing as an exact solution. But we in mathematical programming will push a little further and say that given some specific formulation, we will talk about an exact solution versus heuristics and there are people in the other room talking presently about heuristics.

It seems to me that before we can even talk about exact solution methods, we have to have perhaps one, and only one, well specified objective function. Otherwise, we cannot talk about the exact solution. And it seems that in some aspects of some of the problems discussed yesterday, there was no single agreed upon objective function. Nevertheless, there were heuristics being applied, by Carroll Keyfauver, which were producing, in some sense, feasible solutions. They could not be evaluated very well because there was not a single objective function to evaluate them with. If there were several criterion functions, it would also be difficult to evaluate them against other possible solutions.

So we might want to start thinking in terms of problems with a single objective function. We might think later about how to generalize that.

CHRISTOFIDES: Yes, I absolutely agree with that. In fact, in Figure 12.2 I have a point number 1 with some relevant items. I must admit that some of the problems discussed yesterday seem vague; I do not have a feeling that I know such a problem in a strictly defined sense. Although one can have a vague, general problem and apply a heuristic to it, one cannot do that with an exact procedure. One cannot apply an exact procedure to a vague problem.

EXACT APPROACHES

1. DEFINITION OF SUBPROBLEMS

- © ABSTRACTIONS & GOOD PROBLEM PARTITION ARE CATALYSTS FOR THE DEVELOPMENT OF EXACT ALGORITHMS.
- © WHAT ARE THE BASIC SUBPROBLEMS?

2. EXACT METHODOLOGY

REPORT EXPERIENCE WITH:

- © BRANCH AND BOUND
- © BOUNDS - { - LAGRANGEAN RELAXATION
 - { - STATE-SPACE RELAXATION
- © BRANCHING
- © BENDERS DECOMPOSITION

3. EXPERIENCE

HAS ANY REAL PROBLEM IN SHIPPING EVER BEEN SOLVED EXACTLY?

- © EXPERIENCE TO REPORT?

4. FUTURE

WILL EXACT ALGORITHMS BE ABLE TO SOLVE REASONABLE-SIZE PROBLEMS IN THE NEXT DECADE?

FIGURE 12.2

So I think that there is room here for a very formal description of a fundamental subproblem or a set of fundamental subproblems that are of interest in the routing and scheduling of ships. There is an advantage that road vehicle routing has over shipping in that in the late 1950s Dantzig and Ramser wrote a paper which said: "This is the vehicle routing problem. Your problem may not be exactly this, but this is what we call the vehicle routing problem." And that is the core problem which has been examined by many people. The reason that the mathematical state of the art in road vehicle routing is perhaps more advanced than that in ship routing and scheduling is that we have had the benefit of somebody giving a succinct description of the core problem at an early point.

There are many problems to which mathematical methods are applied that have many complications, but the catalyst for the work that has been done in that area of road vehicle scheduling has been this strict definition of the core problem.

MACLEAN: I think that is really right to the root of this matter. I think it became fairly obvious yesterday, and it has certainly been elaborated on in the sessions that I have been in today, that the problem has not yet been defined. Everybody is working on some piece of the problem. The input and the output are very narrowly defined with respect to specific operational requirements. But we do not have a definition of the overall problem. We do not have a definition of the objective function. Thus, we do not have a specific capability to evaluate the performance of any of the approaches.

GRIER: One thing I would like to say along that line is, the government's objective is to achieve the minimum cost and use the minimum assets.

MACLEAN: Those are two different things. Minimum cost is one thing and that may imply a short-term decision. Minimum assets has long-term implications, and how these two relate has not yet been identified.

GREER: That is right. But you can relate them in dollars if you set dollars as the objective. No matter what you do with the subproblems, you do not really gain your objective in the government until you lower the number of dollars that you spend in order to do the job that you have to do.

MCLEAN: Okay. But is it now or later?

GREER: It is over some visible horizon--a year, two years, whatever you can look at. On the commercial side, it is not the same. The objective, of course, is to maximize profit, which is rather the opposite.

MCLEAN: Perhaps you want to maximize return on assets instead of just profit.

GREER: Yes, perhaps return on assets. I work on the government side, so every time somebody says he has a good solution to that problem on that route or a couple of subroutes, I say: "That is fine, but how many ships do we get rid of? How much money do we save at the end of the year?" Well, in fact the ship is sitting over here rather than sailing over there. That would not save anything.

JAIKUMAR: You spoke before of Dantzig and Ramser. Dantzig and Fulkerson defined the tanker scheduling problem in 1954, but only about three other people picked up on the application and did some work on it. But the 1954 paper instead was the seed for all the network formulations. It is interesting that in the Dantzig and Ramser paper there was a catalyst which looked at the problem in road vehicle routing and people focused on the application: in tankers, it went to a different level of abstraction and people developed algorithms and the whole network type of analysis. So perhaps even if there is a catalyst, it is not sure what it will generate.

ENGLISH: I would like to say a few words about why I think we are gathered here, and I will talk from the Defense Department point of view.

In Defense, we have two problems. One is to run an efficient, economic peacetime steamship company. And I think that Bob Greer was talking to this. The problem involves long-term capital decisions about how big a fleet should be.

We have another and more immediate problem, of more concern to us as citizens and taxpayers, which I think actually prompted us to assemble. It is that we are about to embark upon building fighting models in case the country has to go to war. How can we plan best to do it and if we must do it, how do we do it right?

I think the biggest immediate expenditure is going to come in connection with these models. In terms of anything that may come out of this symposium, I think it is possible that that may be the more immediate thing that any of us who choose to talk further after we leave here will be talking about.

And in terms of who we all are, to use a medical analogy, I think we are fortunate to have absolute specialists here, because without you guys who have actually wielded the scalpel, practically, this symposium is pointless. Then ranging down from the real pros, the guys who do it as a part of their livelihood, there are others of us who are still in business, the analytical business, but we are rather the GPs. We have never really done it to the extent that the pros have done it. And then we have lay members.

And with regard to the GPs and the lay members, I think, to some extent, what we are trying to tell the pros, the surgeons, the real specialists, is just where it hurts.

Because when we leave here, we are still going to need your help. I do not know quite how we are going to do this. With the forbearance of the specialists, we may try to assemble an informal panel of people to help us with this because it involves millions of dollars. It is a matter of lives in the future. It is now a matter of national security. It is important. We may try to retain our connections with you specialists through a panel arrangement where we can at least phone those of you who do have an interest, financial or intellectual or whatever, in the problem.

I just felt I wanted to say that so that whatever else happens here, at least we realize that this was the reason we provoked this kind of meeting.

CHRISTOFIDES: Perhaps a way of doing this, if you are more accustomed to modelling, is to come out and say to the specialists, here is a model, more than one model. Does this look to you as if it is interesting and useful? Someone says no. Okay, here is model 2. Do you think this is useful? Perhaps this is a better approach than trying to define the problem *a priori*.

ENGLISH: It could be, and let me mention another private agony about this. As you know, because of the nature of what we are doing, there is a classification problem. I think we have all noticed there is a time constraint, and there is also a security constraint. And sometimes, when we are trying in this forum to tell each other about things, these constraints get in the way; we can talk in broad terms about the full picture but we are constrained in what we say by time and by security. But the point is, as we begin to know one another better, I think that we can begin to understand how best to go at the whole problem.

MACLEAN: One of the things that we were just talking about in the other room is the fact that in order to be able to attack problems such as this, and use the experience that has already been exposed, we should realize that the solution is not going to be ready tomorrow. It probably will not be ready for any foreseeable time less than two to three years, maybe even five.

And in order to be able to even get together a working effort of consequence, there is probably going to have to be a sizable amount of aggregation. And I think that if you look at what Bruce Bishop was talking about earlier today, if you talk about the real world as it is presently operating commercially, there is aggregation there. World-scale is a good example of aggregation. They have broken the world up into a whole series of areas so that if you route a vessel into an area it does not really make much difference what port you go to.

When the size of the military problem was put up on the board yesterday and elaborated on today, it boggles the mind as to how computational capabilities can be brought to bear on it. It is beyond anything that has presently been done. And it seems that you are going to have to reduce the size of the problem in order to be able to get solutions to it in the foreseeable future.

I would suggest that these are things that have come to my recognition, and I ask if they agree with the observations of others?

HASKIN: Let me go back to your definition of a subproblem, which is somewhat related here, and what Tony English mentioned before about putting uncertainty into the problem. Have there been solution approaches where you have a practical subproblem, which might be the actual routing for a given period of time, and you take it out of the model to the real environment, insert the factors, and then go on with the model? Have there been models written that develop that way?

CHRISTOFIDES: Yes, I know one model for the collection of milk from farms where the amount of milk you have to collect is uncomputable. It depends on the rainfall of the period, how the cows feel, and so on. That is a model which has incorporated in it uncertainty in terms of the quantities and in terms of the places you have to visit. There are many cases; this is just one. Certainly, in fleet routing and scheduling, there are many, many cases of a model which was originally produced for deterministic purposes being extended to a stochastic model.

I put a couple of points here in Figure 12.2 about methodologies that perhaps can be useful in dealing with problems like the ones we have been talking about. Really, the reason I put those points there is to see if anybody has any experience to report in using these methods. Yesterday Jai Jaikumar reported some experience. Does anybody else have any experience with using any of these techniques in problems that may be useful in ship routing and scheduling?

BAKER: There was a small steamer operation problem to which I applied a branch-and-bound scheme. I talked with the people at the company, and they said in their particular planning problem time considerations were of major importance to them because of particular delivery time windows

and depleting inventories; they were assuming a constant depletion rate over time. Most of the factors involved in navigation would affect the time considerations to some extent.

So I formulated a model based upon minimization of total lead time (or voyage time), and basically came down to using a standard branch-and-bound scheme. It was a very small problem--three tankers and 25 ports. Solutions were obtained in less than half a minute of CPU time, rather quickly.

CHRISTOFIDES: What bound was used in this scheme?

BAKER: I was actually using a dual longest path problem. Because of the rather unique nature of the formulation, one of the relaxations turned out to be a longest path problem, as in a PERT Problem. That was the bound that was used.

SOLAND: Was it a Lagrangean relaxation?

BAKER: Actually, I did not have Lagrange multipliers. I was able to just relax constraints. Each subproblem of the branch-and-bound tree turned out to be network problem.

SOLAND: Did you relax the problem by throwing away constraints or by relaxing integer constraints?

BAKER: By relaxing integer constraints.

CHRISTOFIDES: I had some experience with branch-and-bound using state space relaxations for bounds. This involves formulating the problem as a dynamic programming problem originally, which is very easy to do for routing problems, and then relaxing the state space into another lower dimensional state space in such a way that one gets a lower bound by solving the relaxed recursion, the recursion in the relaxed state space.

One can get some very good bounds for routing problems this way. In fact, this branch-and-bound code with the relaxation was the one that was used in the shipping problem that I mentioned earlier; again, a small problem. It involved only six ships delivering refined oil products to about 45 ports. In that case, the problem was solved within a fifth of one percent in a completed search taking 10 seconds.

AD-A123 464

PROCEEDINGS OF A SYMPOSIUM ON CARGO SHIP ROUTING AND
SCHEDULING HELD WASH. (U) GEORGE WASHINGTON UNIV
WASHINGTON DC INST FOR MANAGEMENT SCI. R M SOLAND

4/4

UNCLMRSSIFIED

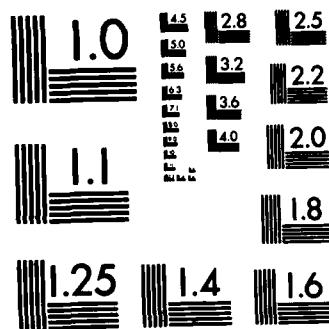
15 DEC 82 SERIAL-TM-69250 N00014-82-G-0016 F/G 15/5

NL

END

FIMED

DTC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

GREER: We ran into a similar problem in terms of time. We ran six dry cargo ships. They were running east coast, gulf coast, California, the Far East, and all of the Pacific. There were about 25 or 30 ports. We ran the schedules over and over again, trying to determine how to route those ships. And you are correct. Time is the important thing. Then you find out that if you get your time down, then you get more cargo. All of a sudden it generates itself, if you get to the point of reducing your time.

CHRISTOFIDES: There was a problem we had that was not mentioned yesterday. Maybe it was unusual. It was the queuing of ships at port. It was one of the most important considerations in that a lot of time was spent with the ship in the port doing nothing because there were ships in front of it unloading. As part of the optimization, one had to consider the expected queues at different ports with different times of arrival. And the actual call time, the time to unload at the port, depended on the actual time you arrived at the port, simply because it was an expectation we used.

ENGLISH: That was part of the Vietnam experience, and I think that everyone is hoping that we will never have another one like that. I think any of you who have familiarity with the Vietnam experience recall that one of the nightmares of Vietnam was that there were ships waiting to unload for weeks, and maybe even months.

BISHOP: Everybody is saying that they have a small shipping problem, with six ships and 45 ports. That is probably the typical problem. They do not get much larger than that. You can segregate the interactions down to either a small number of ships and a large number of ports or a large number of ships and a very small number of ports. And you do not have a lot of routes. You do not have a lot of streets that you have to deal with; you have one great circle road.

So it is not like the land models. There you have lots of trucks and lots of products and lots of different ways to get there. In shipping we can generally break things down into small units, which cuts down the aggregation considerably.

MARSTEN: We have to be careful. There are many different problems that are being talked about here. What we were talking about in the liner operations, I think, is totally different than what MSC is interested in, with potentially thousands of ships.

BISHOP: Yes, but what I am saying is if you have got a situation where you have a lot of ships, you are probably trying to move a lot of stuff into one area. That is the military problem. Then you have a limited number of ports to deal with. There are only 2000 ports in the world that can handle cargo, at least for tankers.

ENGLISH: That is true. For any given military situation, it does narrow it down.

BISHOP: I mean, on the west coast, you have got perhaps three ports you are going to use. And in the case of Vietnam, you have perhaps four that you were going to.

ENGLISH: So it is true that it may be we are frightening ourselves unduly, that we could at least get less shaken by it if we talked in some more detail.

KASKIN: The largest of the problems, however, do get complex because we are talking about using all of those ships. Those are not all used for a single theatre of war. They are used in the case of multiple theatres of war, world-wide, and then you will be using all the U.S. ports and as many ports as necessary overseas. So that is what we have to plan for. The rest of the problems, the smaller ones, will certainly be much simpler.

HALL: Yes. But perhaps, as a way of getting started, we ought to focus on some of the likely scenarios. I think most of us would agree that the most likely scenarios are the smaller ones. And perhaps we can figure out how to do that well, and then generalize.

KASKIN: Nobody has yet said to us as operators what is most likely, as a policy decision.

FOARD: I do not think you can find anybody in town that will identify what the most likely scenario is.

MARSTEN: Responding to the request for some computational experience, I have worked on air cargo problems. We have a problem with 50 cities, each of which has cargo for the other 49 cities. I built a mixed integer programming model where the users put in some network to be used in connecting the cities and then the model decides how the cargo is to be routed through the network as well as what type of aircraft (from some list of, say, half a dozen different alternative types) is to be used on each part of the network.

We have gotten very good results by using branch-and-bound, basically by solving a linear program and then using some clever heuristic to round. The aircraft theory used is based on some knowledge of actual aircraft. It is not a simple rounding, but a clever rounding.

That problem seems to me a lot easier than the problems dealing with ships, particularly because, as was said the other day, the days are separate and each night is a buffer. In our case, the daytime is a buffer between separate problems and you do not have a problem of where the planes are going to be tomorrow. You do not have accumulating uncertainty.

CHRISTOFIDES: From what you were saying, it seems that nobody has had any experience with either decomposition methods or cutting plane methods for solving these mixed integer programs. So is everybody working on branch-and-bound?

RONEN: It seems that way.

CHRISTOFIDES: Now I see that we are running out of time. Let us just go the last issue in Figure 12.2. Do you think that in the next 10 years, say, we can solve exactly some shipping models of reasonable size, problems that will be useful to the shipping community? Do you think that the methodology exists, or that it does not and will be developed?

BAKER: I would like to make one comment about the application to a tramp steamer operation. In the very first minute that I was associated with the company, the profit motive was firmly established as the objective. It was clear that everyone who was working with the problem wanted to know the profitability of a particular voyage.

And when you analyzed it more closely, it was determined that the leasing charges for the tanker were a function of time, the length of the voyage, the administrative charges, the crew charges. Port charges could be added as you went along. But almost all the major costs associated with the operation of the tanker had something to do with the time factor.

So we made an assumption that that is the direction in which the objective function should go. As I think we all know, once you have a clear statement of the objective and can mathematically compute the constraints, then you can apply exact algorithms.

BALLOU: How about that as far as the other transportation modes? Has that already been done for trucks; has it already been done for rail?

CHRISTOFIDES: In vehicle routing, if you think about the next decade, I think it will be quite possible to solve most medium-size problems exactly. The difference between the state-of-the-art now and the state-of-the-art four years ago is considerable. Average size vehicle routing problems involve something like 150 calls per day and 15 trucks. That is from a single depot. That kind of problem could only be solved heuristically and without any degree of guarantee four years ago, the errors being perhaps anything up to 10 percent. Such a problem can certainly now be solved to within 1 percent, and in some cases, to within a tenth of a percent.

The rate at which the state-of-the-art with respect to that type of problem has improved leads me to think that in the next 10 years we will be able to resolve it exactly. I am talking about the abstraction which is the core problem of a practical problem. I am not talking about practical problems with lots of different complications, some of which are subjective, some of which cannot even be objectively stated. Those are always going to be difficult. They are not mathematical problems.

SOLAND: I will venture a guess and say yes to your question, in the following sense. I think that people are going to cut into the large encompassing problems enough to specify an abstraction that is significant enough to work with. It is going to be a suboptimization in the larger sense, but it is going to do a certain amount of scheduling. It is going to be a very important component of some overall system. The

problem that results even from the abstraction is going to be big enough to be significant, something on the order of the problem that Jai was talking about yesterday, or bigger. But the techniques that are available and becoming improved and more available are going to be able to cut that down so that with some well-formed objective function, which may still be somewhat difficult to relate to the larger problem and system, but still one that is agreeable to work with, one will be able to optimize the problem to within, say, one or two percent, and I consider that exact. I do not think that anybody is going to ask for more than that in terms of uncertainty of data and uncertainty of the suboptimization within the larger system.

MACLEAN: In view of the fact that it took about 4-1/2 years to bring the APL container optimization problem to its present state, in view of the time it took to bring the Airco work to its present state, would anybody venture a guess as to how long it is going to take to overcome the problems that an exact algorithm must deal with for the large defense problems as stated? With the present state-of-the-art, are we talking about 6 months or 3 years?

PSARAFITIS: Are you talking about full-scale implementation?

MACLEAN: Yes, presumably so. Even if it is a trivial part, you are talking about putting it into place and trying to validate the thing and have it as an operational tool.

KASKIN: It is a function of the amount of money you want to put into the effort. If you get all these people here in this room on a Manhattan-type project, you will get some results rather quickly.

MACLEAN: Do you think 6 months would produce results?

RONEN: Probably. If they are ready to put up the money.

ENGLISH: Obviously, there is another side to the question, because we do not have the money. Although we have more money than I have seen in some time for this sort of thing, we do not have all the money in the world. There

is the other side to the question that has to do with the other counterpart of this session, and that is that people have done some very persuasive work using heuristics. And obviously, one of the reasons we are here is to try to determine the best way to go next, whether it is one or the other.

But obviously, since we do not have all the money in the world, and we do not have all the time in the world either, there are going to be some very tough management decisions, because we do not know when we are going to be invited to another war. So we want to be ready, as ready as we can be.

So there obviously is a trade-off in this regard. And to try to get some better understanding to approach those management decisions is why we are here.

GREER: There are two problems. There is the dry cargo problem and there is the petroleum problem. The petroleum problem was worked out fairly well down at Texas A&M. They got within one ship or two ships of doing the problem as well as the people around the tables did by moving the boards around, etc. Now if you can operate that way in a peacetime environment and then you can shift into a wartime environment, you have a model that will continue to work for you, perhaps one that is good enough in wartime if you get within one or two ships. In reality, you do not know what the master of one ship is going to do compared to the master of a different ship. You do not know that well and the model does not know either.

I think that is workable. From what I have seen and heard I think that is definitely workable. I think when you get to the dry cargo, of course, you have a more complicated problem. There are more different kinds of ships and different kinds of cargoes. There are more ports perhaps. Someone said 400 ports, but I do not think that Defense deals with that many. We go to perhaps 200 ports for dry cargo.

But there again, one would hope you could build the model and do the same thing, i.e., have a model that gets fairly close in peacetime

and also lets you shift into a wartime situation. MSC has the problem of knowing the world, so it can draw, as has been described before, on all of the various resources to build up the fleet to a large size. But you must have a model that can handle it. I think it can be done. The only problem that you are going to have is whether on a day-to-day basis in peacetime people will use it, because they will be able to beat the model a little bit in peacetime. And then when you get to where you need the model, they will probably be willing to let that model operate. They will say, "Oh, that is great. I throw my hands up now because now it is too big a problem to deal with quickly." They will be glad to have the model.

PSARAFITIS: I do not think it is the number of ports you should consider as far as complexity is concerned. I think there is a big difference between a ship dealing with a homogenous commodity like oil and a ship in a dry market where you get hundreds of different commodities to be carried on different ships or on the same ship.

So I think it might be misleading to draw conclusions from models that have been applied to the tanker market.

GREER: Right.

ENGLISH: I think that that allows me to make a point that came to mind when Bob Greer was speaking. He said "the" model. I am at some pains to explain that there are different things and different priorities that we are interested in. There is the peacetime Defense Department steamship company operation, which at one time was very important when there was not such a war-fighting interest. And then there are these war-fighting considerations.

Discussion Group 6 - Heuristic Solution Techniques

Leader: John J. Jarvis

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA

JARVIS: I would like to discuss the term "heuristics" that has been bandied around quite a bit in some of the other sessions. I would like to see if we can assess the state of the art in terms of the development and the application of heuristics.

I want to try to bring those people who are developing heuristics together with the people who are applying the heuristics, to see if we cannot get a greater understanding of the needs and the capability, and maybe from that get some feeling for what lies ahead, what the possibilities are.

I would just as soon stick to the possibilities and not the impossibilities. I guess we need to start out by trying to define the term "heuristics." I am afraid we have to put a modifier on that, because we can have a heuristic technique. You can give me a model, and my technique is a heuristic technique if it does not guarantee the optimal solution to your model. That begs the question, because your model may not be an exact representation of your problem.

So I would like to view it in a broader context. I would like to talk about a "heuristic approach," which is a combination of giving up some of the constraints and relaxing the modeling process, as well as relaxing the solution process, so that you do not get an exact answer to the problem that you tried to solve. In that sense, I want to speak of heuristics.

DOUGLAS: Why can't you get an answer with the understanding that it may or may not be exact, but it is not necessarily 100 percent appropriate, or something like that? Maybe that says the same thing.

JARVIS: The primary reason that you cannot get exact answers to real problems, I believe, is that, one, the problems are usually so complex that they cannot be even fully modeled with some mathematical modeling technique. You cannot step in out of the social, political, economic

constraints to get a real model in the process. There are too many factors. And the second thing is if you are able to get that kind of model, you cannot solve it exactly--when you start to deal with the computer--in any finite time.

So I think there are two reasons why exact models are so difficult to develop for real problems. Consider the problem Jai talked about. They take a problem and develop some kind of model that, in theory, has a large number of columns and a large number of routes, but then they limit the number of routes that they generate, so they have not fully modeled the problem.

PENTIMONTI: I would comment that maybe the effectiveness of that kind of model is limited by the ability to do it exactly. It may be able to be done from a more cost-effective point of view.

JARVIS: Let me run through some of these ideas to put the whole thing in perspective, try to set the discussion, and then I want to open it up. (See Figure 12.1.)

This is what I would like to get some feeling for from the group: Can we each fit ourselves into this matrix someplace in terms of where we are personally working? What I would like to do is get each of you to say where you see yourself within this kind of matrix, either as a developer or an applier of these things. And I think I would like to try one cut at categorizing heuristics, those that just guarantee feasibility.

I just have to have a solution. I have to have one today. I do not care if it is the first solution I get, because of the time it takes, or the people I have to allocate, or the computer availability, or whatever; I just want a feasible solution. I will keep working until I get one.

From there we try to improve the quality of the solution. One of the first steps is to say "well, let me take some rules of thumb that have either been generated by previous computer experiments, personal knowledge, or watching the schedulers and trying to abstract

HEURISTICS

Group

Goal: Assess the state-of-the-art of DEVELOPMENT AND APPLICATION of heuristics in ship scheduling

Method

Types:

- Feasibility
- Sensitivity Rule-of-Thumb
- Geometric based
- Optimality based

	Manual	Interactive	Automated
	X		

Quality

- Theoretical analysis
- Experimental analysis
 - improving practice
 - standard test problems

FIGURE 12.1

what they are doing--trying to emulate the schedulers." I can develop some rules of thumb by maybe allocating the next ship to the closest port, or by taking the one that has a minimum number of delay days until the next allocation, or on some other basis.

Then there are those heuristics that are geometry-based; I would not put them in any order of quality. They have to do with the geometric aspects of the shipping problem, considering the locations, the ports, and where you position the ship next in terms of having the necessary capacity in terms of a look ahead to the next geometric location. And then there are those heuristics that are optimality-based, where you write down an optimization model.

I am not saying that this is "the" particular model that solves the problem, but it is a model that approaches the solution of the problem. Then you back off from that. You start doing some relaxation, like Jai talked about. They wrote down a model and then they treated it elastically. They got rid of some of the hard constraints, and made the things react a little bit more simply, and then they could start playing with that. And you might even relax even more, to where you would get down into some penalty structures, and then from that you could have very simple kinds of traveling salesmen approaches, or other submodels that would help you identify good routes that you could then start putting back together.

One of the terms that has come up quite a bit with regard to assessing heuristics has been their quality. One of the things that has not been talked about too much, and really, there is not too great a state of the art, is in terms of the theoretical analysis. No matter what problem you give me, I can always get within 100 percent of the solution in terms of the measure of effectiveness. Typically, there is not much said about this in practice, because these bounds, this assessment of quality, is not very good. In those kinds of problems you have to take into account all of the possibilities that ever come up in anybody's lifetime.

So you have to be very loose about that. What are concentrated on are the two kinds of experimental analyses. One is in terms of measuring the quality of heuristics, or any method, in terms of its ability to improve practice. That is, we put this together in our shop, and we find some kind of improvement in terms of movement, you know, the amount to be moved, or delay-days, or a profit-and-loss statement, whatever your measures of effectiveness might be.

I do not say "measure," I say "measures." I do not think that there has been enough done about trying to develop some standard test problems, that is, some ability to start comparing my method to your method and start having some cross-fertilization of ideas and methods. There does exist a set of standard test problems in the area of vehicle routing and delivery.

There is no reason why we cannot begin to make available some standard test problems in ship scheduling so that we can begin to compare and cross-compare our methods with other methods on some kind of common basis. I think we need to think a little bit about that, how to begin to generate that.

There are two other things I want to throw out for your consideration. (See Figure 12.2.) One has to do with problem complexity. Now we are getting on the applications side of it. What kinds of considerations with regard to the problem force us to move away from, say, optimality towards just feasibility, in terms of selection of heuristic techniques? I think there are a number of things having to do with problem size, certainly the number of ports, the number of commodities, fleet kinds of constraints, size, and type. We have port constraints, temporal considerations, time windows, number of time periods, degree of discretizing the problem.

One important aspect of the problem, one which would cause you to move away from manual to interactive is the stability of the scheduling environment.

Problem

Complexity:

- Number of ports
loading, discharge, etc.
- Number of commodities
- Fleet constraints
size, type, etc.
- Port constraints
- Temporal considerations
- Stability of scheduling environment
- Interface with other systems

Choice

Considerations

- Availability
- Ability to react
- Operator control
- Quality
- Understandability

FIGURE 12.2

Suppose we have a person doing it manually. He has been here for 20 years, he has done it a lot, and he knows pretty much what is going to happen. If it is in a stable environment, I think that is probably true. If you move to a more unstable environment, or start looking at a variety of scenarios, it gets much harder for the human to still gain control of what things were good and what things were not good. You have to start relying on greater and greater amounts of information. As soon as you get into the information processing, then the human starts to get weak, and I think the computer plays an important part.

There is also the interface with other systems. If I not only have to solve my problem, but I also have to worry about how the units are getting to the ports and when and so on, and I start interfacing with someone else, either by passing information up to someone who is coordinating the whole process, or just talking to the people who are doing it and trying to get a match at the boundaries, the more complicated the whole system becomes.

There are a number of choice considerations with respect to heuristics. One of the most important choices is whether a method is available. Quite often we cannot spend the time and effort to develop a whole brand new system. Maybe we can move into it over time, or maybe we begin to pull in pieces here and pieces there of different kinds of systems to build scheduling algorithms.

Another important choice consideration is the ability of the method to react--to react to changes in the environment due to scheduling modifications.

Another important aspect is operator control; whether or not I am obligated to accept it and whether or not I can interact with it. Can I at least get enough out of it to develop alternative solutions so as to deal with some of the nonquantifiable constraints and considerations that I could not build into my model? I do not want just one answer. I want a range of choices to give me some flexibility. And then certainly there is the "quality" of the output, whether or not it tends to compare well, based on some measure that I have to tell me whether I am doing a good job.

I do not want to take any more of your time to just sit here and talk about these things blindly. What I would like to do is open up the discussion. Maybe we can have some people react by giving us a feeling for where you think you are relative to this kind of matrix. Maybe you can talk a little bit about what you are doing and fit it into this framework.

Tom Baker, why don't you tell us a little bit about what you think EXXON is doing?

BAKER: Our interactive scheduling systems are in the bottom center, optimality-based interactive systems. By "interactive" on that bottom line, we mean really the whole bottom row, because it is up to the user as to whether he wants to operate more in the manual mode or the automated mode. But that is what we mean by "interactive." We blur that distinction, I guess, or definition of "interactive."

JARVIS: What is your definition of "interactive"?

BAKER: The user can come into the system and he can do it optimally by letting the computer schedule the whole thing by itself, or he can make every decision and work his way completely through it, using the optimal solution, or he can use different parts of the optimum. That is what we mean.

DOENGES: But when he is operating by hand, is he still striving for optimality, or just feasibility?

BAKER: Some of the infeasibilities are flagged on an exception basis so he can see those. In general, we can find an overall objective function that he is always working against, and he sees that on all of the screens on which he is making decisions. There is a penalty cost for not respecting due dates, and so forth, built into that overall objective function, so with any decision he makes he sees the effect on the overall schedule and so does the heuristic optimizing model. It is somewhere on that bottom line.

JARVIS: I do not mean to imply that down is better. There are many very complex problems which only admit to feasibility. I would be interested in hearing somebody talk about that.

MENTZ: Would not most everyone operate on a diagonal from the upper left to the lower right?

JARVIS: In what sense?

BAKER: I think the answer is no.

WEBSTER: You do not want to go to the fully automated in general. It is a possibility, but you would throw away the user's whole experience if you aimed for that. You have to have the ability to tap in a little bit. The center is interactive, and the right-hand side is automated. By that, I presume you mean fully automated, so that Tom pushes a button, and calls opt, and then it does everything. I think that would not be my view of your goal.

MENTZ: I guess what I meant was that I am not sure that I understand all of the elements in the four by three matrix there. I do not think there is a reality to an optimization routine being done manually. I am saying, if you are down near the bottom of the matrix, you are going to be pushed to the right. And if you are up near the top of the matrix, the right does not make any sense. So, you are operating on the left.

WEBSTER: I agree with that.

MENTZ: I am saying there is only a diagonal to deal with. Then the question is where is the best place to be on the diagonal, and I think it is somewhere in the middle, toward the upper left-hand corner. And you have a distribution along the diagonal.

WEBSTER: The only way I was disagreeing with you was with respect to how broad you consider the diagonal. One by two is okay.

MENTZ: I do not know the answer to that question.

DOENGES: I would not agree that the upper right does not make sense, especially for MSC. I think maybe they are not or should not get hung up on having an optimal schedule. I think they are looking for good schedules, feasible schedules. They might use an optimization technique in deriving them, but the feasibility is the key thing. And certainly you might have a very highly automated system to define what is feasible.

JAKOWSKI: But the question is this: if you come up with an infeasible solution, is it because you have not put in the proper work to come up with a feasible solution?

DOENGES: Your model may say it is not feasible because of the limitations of the model. That is a possibility and an advantage to optimization modeling in that it shows you the direction to move in so as to resolve the infeasibilities.

I do not have a lot of heuristics experience. Are there heuristic methods that can also show you the direction in which to move?

JARVIS: In fact, the idea of a heuristic or any kind of method should be to move you either towards feasibility or towards optimality.

DOUGLAS: Even if it is doing just two solutions in a row and finding which is better?

DOENGES: I think it may be the same thing. But let me put it another way; if your model, whatever its type, determines that the solution is infeasible, what kind of information is readily available from the output of the model to show you what would need to be changed to make it feasible?

JARVIS: That is a good question.

DOENGES: I guess I am suggesting optimization approaches are better than heuristic approaches in the sense that they do provide that information.

JARVIS: Checks on infeasibility, doing well when it tells you it is infeasible, there is lots of marginal value information, for example, that lets you know what to change to make it feasible.

BRIERRE: If you take a penalty function approach for the infeasibilities, between the top line and bottom line, there is not much difference.

DOUGLAS: And then there is not a great need for the right-hand column-- if you agree that you need people in there. If people are in there, then it is going to be interactive to some extent. So you are sort of narrowing the thing somewhat.

COPPINS: If you mean on-line interactive, that is different from saying that if there is always a person, there is always a feedback loop. I do not think anybody believes that you can push a button and get an answer that everybody is happy with. I think even in Tom Baker's case, when he says they just push a button, what happens is they then look at the results and say "Do I really like that?" It is always interactive in that sense.

JARVIS: Do we then begin to believe that there are only a couple of blocks left?

DOUGLAS: It is starting to look that way.

BAKER: I accepted the bottom left because of the use of the overall objective function. The user is basically forced to define his overall objective function. So, even if he is doing manual scheduling, he is weighing how he does against the objective function.

PENTIMONTI: Using the model as a computer?

BAKER: Like a cost-function evaluator that evaluates every decision.

PENTIMONTI: It is not an optimality-based solution.

BAKER: When you get right down to it, optimality is a luxury of academics. This cost-function evaluator we assume to be the best model of the overall cost and inventories that is in the computer. But the

model that is in the human's head is an even better one.

PENTIMONTI: That is my point.

BAKER: The other models were used to try to optimize, to go back and check against this simulation model.

PENTIMONTI: That is my point. We do the same basic thing. I would put us up on the top line.

BAKER: If you define optimality according to the overall objective function that he has in his head, then you can do manual optimization.

JARVIS: I am interested; did you say you would come across the top line in your operations?

PENTIMONTI: As I mentioned earlier this morning, we do not have any objective functions within our scheduling simulator. We simply use it for feasibility checks, and then, of course, use it as a computer to tell us what the returns will be, profits, or some other indication of financial return, on that schedule plan. We are on the top line all the way, maybe the second one; some definition of either just feasibility or both feasibility and data generation.

MENTZ: If the bottom line of the ledger was construed not to be a very good result by someone, saying we ought to earn more on that voyage with that cargo delivery sequence, then that would be using a rule of thumb to evaluate that solution?

PENTIMONTI: That is right.

WEBSTER: I do not know where it fits in exactly in your matrix, but I think, to the user, maybe one of the important things is the marginal situations, that is, the sensitivity to various pieces of the input data, like the model not quite being correct. Certainly the input data are not quite correct either. And then the user is most frightened about what will happen if cargo does not show up, or if the ship breaks down, or the propeller has fallen off, etc. There are all sorts of terrible

problems that can arise.

Optimality from the point of view of the user may be viewed by him as his risk or his exposure more than it is used by him as totally just a bottom line. That is a consideration that is sometimes hard to quantify.

JARVIS: Which way do you think we are headed here? Do you think we are headed towards the middle or toward the right, towards the bottom? Where are you personally headed?

MENTZ: I think we are on a diagonal from the upper left to the lower right. We are somewhere in the middle, and we are struggling down toward the lower right. But we are afraid of going too far. We are afraid of the unknowns in that direction.

DOUGLAS: We are not afraid of going too far, but just realizing that you should not go too far.

MENTZ: Realizing that you should not go too far?

PENTIMONTI: You cannot do that economically.

DOUGLAS: I am saying it is not good to go too far.

MENTZ: The comment that was made by Tom Baker of EXXON to me sums up the entire day and a half so far on this question of optimization in marine operations. The point is that optimization is a luxury for academics. I hope that will be in the transcript, because I think to me that sums up the day and a half that we have had so far.

JARVIS: I do not think we will have any problems in this room. I think that is a problem in the other room.

MENTZ: Through a totally different set of experiences, I have come to the same conclusion.

JARVIS: On the other hand, I do not think we can throw out modeling totally and say we are going to have our people out there doing it.

MENTZ: I do not think we said that.

BAKER: I, too, do not think we said that.

JARVIS: We are trying to improve the capability of the techniques to do a better and better job of doing what we think is good. When we create an objective function in a model, we are trying to reflect what we think is good.

PENTIMONTI: From a liner's point of view, Bill Webster hit a note that was resonant. What we like to do, once we have set manual schedules and used the computer to check feasibility and do computations, is to check sensitivities. We will run the exact same model with a different port time and a different amount of cargo, and high and low estimates on revenue and so forth. And, thus, we have a tool we can interact with and yet still come up with what we consider, based on everything we know, rational and optimal solutions. It is totally interactive, and it starts with a very simplistic feasibility type of simulation.

DOUGLAS: It may use optimization features. If you have to evaluate different solutions coming out and they are triggered by or they are calculated by optimizing procedures, then you have the best of both worlds.

PENTIMONTI: Precisely. But the model itself does not do optimizing. The optimizing is only done by the interaction of the people who are using the model, using the tool, and their optimization will really come from their interaction with it, knowing what sensitivities are realistic and checking those sensitivities through the model to come up with ranges of solutions so that one can see a high and a low range on returns based on what are known to be the problems and the reliability of the entire system of variables and data that are used as input to the model.

DOUGLAS: If you are far enough out, I do think you have optimization procedures that can look at that without getting into the scheduling function, which, I agree, does have a problem with using optimization procedures.

MENTZ: Would you consider changing the title for the second row to "Sensitivity?"

JARVIS: From "Rule of Thumb" to "Sensitivity?" (Both appear in Figure 12.1.)

MENTZ: Or adding another row. I think the point that Gene Pentimonti made is very well taken. The modeling efforts, as used in his company and many others, use the capability of the models to look at the various sensitivities and alterations in basic input factors. He might be saying that he is operating primarily in that box, interactive sensitivity analysis.

And maybe with Burnie Douglas' comments, you are reaching out and using some of the other techniques to help you. But basically, you are focusing on interactive sensitivity. Does that make any sense?

PENTIMONTI: Yes, it does.

JARVIS: It characterizes how we are using it, at least. I did not try to indicate what the models would be used for, but how they would be constructed.

MENTZ: Is sensitivity a use?

JARVIS: I would think that would be a use of a model, to conduct a sensitivity analysis.

WEBSTER: In a lot of these heuristic methods, what you do, in fact, is develop a collection of feasible solutions. The fact that a method is heuristic means that you have not necessarily exhausted the feasible solutions, but you have come up with a collection of, one hopes, good feasible solutions.

Now, Tom Baker has some kind of multiple criteria ranking. I do the same sort of thing with the things I do. And that can also include sensitivity analysis, because that is a way of looking at this collection as well.

I think as soon as you are into heuristics, if you just are after one solution, fine. Often you get a collection, and you rank. So it is partially optimal in some sense.

JARVIS: How good are we doing relative to other people? Do you think it is possible for Tom Baker at EXXON to compare his work with what someone is doing at Standard Oil? Do you think that our problems are so unique that the models we are using on our data cannot be compared to the models others are using on their data?

Is it possible for us to learn from each other, for all of us to improve? And if so, can we find out how? Is it possible to find out who has got a good model to do a certain thing, to begin to extract, draw and improve based on this knowledge?

Do we know how to assess the quality of our models? How do you do it now? How do you do it at Bethlehem Steel, Burnie? When you ran these things, you went through some process of trying to provide information for scheduling; right? You were doing this manually, and then you introduced some degree of automation. Moving in that direction, did you get concerned about the quality?

DOUGLAS: No, because it was not an optimized type of thing. It was recreating what the scheduler was doing, giving him the ability to do it fast.

And just as an aside, one of my biggest days was when I peeked in the scheduler's drawer, where he normally kept a sheet of paper that gave, for a particular vessel, the date here and, say, four days up and then two days across, and three days there, all of the dates listed with pencil and paper, that he had been keeping for 20 years. I peeked in his drawer and he did not have that; he had the computer output, the printed output of the schedule, and he had made the changes on that. That was a big day.

So, no, I was not concerned as to the possibility that the quality had gone down--maybe concerned a little bit that maybe the quality could go up. But that was not a concern at that point in time.

JARVIS: Were you thinking about the ability to react with speed?

DOUGLAS: Yes. It did not take him two or three days to reschedule the fleet. He could do that in an hour.

JARVIS: He took the approach that Carroll mentioned in the last session, trying to start out by giving them what they are doing now, just to help them do a better job.

DOUGLAS: It is a big step toward user acceptance.

JARVIS: It introduces some of the information processing capabilities.

PENTIMONTI: We feel that we can get to an absolute, not optimum, but to pick another word, an absolute acceptable schedule that probably does not have very many variations.

Therefore, the sensitivity analyses we have run on the schedules which are selected manually are just giving us relative information, so that we know relatively how the schedules stand. We know we have got the ten best, by the time they are manually listing them out.

The sensitivity analyses give us confidence that we have picked the best one, the most acceptable one.

POLLOCK: One of the interesting experiences I had many, many years ago was to work in a study group that dealt with shipping management. It was really a question of studying what happened to just a single load.

We took one ship, and we literally followed everything that happened to that load from the beginning to where it was delivered. And what we did was meet with a lot of the maritime managements, vice presidents of companies. We wrote down their opinions of what was going on and what was the problem in the shipping industry.

Now, this was not a scheduling problem, but it was a problem where they listed the things that they thought were wrong. And do you know what kinds of things they listed? "The longshoremen are a

bunch of bums," and things like that. We got a real nice collection of things.

We started analyzing that problem, using a computer and trying to understand what happened. As we got the output, the thing that became obvious was that every one of these people, who were managers in the industry, were dead wrong. They did not know what was really wrong; they did not even understand the system.

The thing that was intriguing to me is that, as this information leaked out, their position changed and shifted gradually. As a matter of fact, the group that did this suddenly got attacked by their managers, who said "What you are coming up with is very obvious." Thank God we had recorded what these guys had told us before, because we would have been in a lot of trouble otherwise.

Now, I have a feeling that when we come down to looking at optimal solutions, we are in the same boat. We are hearing a lot of people who are managers talking about what they know is true. And, I would daresay, if there was some expenditure of funds in this area and people looked into it, we would discover that they really do not know what is going on, and that, in point of fact, there are a lot of techniques that could have tremendous payoff for these people.

Well, it has been 20 years since we did this experiment. And, as I say, the reaction I got was so startling. When we finally pulled out what they had told us and said, "Hey, gentlemen, this is what we have. You wrote us letters. We took interviews." They were just astonished that their position had changed.

In the session I had this morning the thing that struck me was the strong feelings of the industry people that there is not much in the way of payoff. I was struck by Russ Stryker's talk, and I thought gee, this goes back to 20 years ago. I wonder if we really have a situation like that. Paul Mentz has been saying there is not much with optimality. The thing that bothers me is that if the problems are that simple, then maybe we ought to be looking harder at some of these problems and trying out the idea. Maybe that guy who was the

head of the liquid oxygen company was on the right track when he said, "Hey, I want to know what optimality is." It was not just feasibility; he wanted to know how well his company was doing, to get as close to the best as possible.

That is now something that we need more of--more of a spirit of trying to find out. Because, if I look at the maritime industry, we have got some problems that are horrendous. These companies are going to be faced with looking at tougher decisions when Uncle Sam stops supporting them.

And I am saying this--it is a hard thing to say to these guys, but maybe they are going to have to sit down and figure out what is optimum, or work more down that diagonal, to the right.

Okay, that is enough. I hope I challenged somebody.

JARVIS: There was a comment earlier by the man from Chevron, Bruce Bishop. I hate to put words in his mouth, but I guess he went in the other direction. He said that they were pretty much using manual methods in their scheduling. They think that they are leaving some money on the table by doing that, but they are not sure how much.

I think that this is a good opportunity for more and more modeling, more and more information processing, to do a better job of finding out.

Maybe we are doing a good job with what we have, and maybe we are not forced to do a better job. What is it costing us not to attempt some improvement? I do not necessarily mean to go all the way to the lower right-hand corner. I personally do not believe that we will ever get there in my lifetime; maybe somebody else will.

I believe it is really important to have the human involved. I think we can gain from that experience . I think the human needs the ability to have models process information for him, so he can then do a better job of controlling the process.

When I mean interactive, I mean total human involvement. I do not mean just getting the output of a run, and then going back to change the input data, to try to get some different run. I mean the ability to react, enforce considerations, and then go back and see what the local reaction to that might be, or the change in some global measure, so you are right in line, involved in the process. Now, I will back off.

We are running out of time, so let us look at the other two possibilities. Let us get to the application side of things.

Where are we bogging down with regard to the application of heuristics?

What seems to be the biggest problem, in terms of problem complexity, that we are bogging down in with respect to the heuristics themselves?

WEBSTER: I would turn that another way. I would say that frequently heuristics are insensitive to complexity. Double the complexity and you double the time. You do not have so much sensitivity in that direction. I think that is the difference between using heuristics and getting fully optimal solutions; if you double the complexity you may go from the ability to get an optimal solution to a place where you cannot even write the algorithm.

So, if anything, I think increased complexity forces heuristics.

JARVIS: Let us assume we have given up optimality. We are now very happy with heuristics.

Is there not a possibility that the heuristics themselves can be torn up? As Tom Baker indicated earlier, some slight changes can tear up one heuristic but not another one. For example, suppose your time windows get a lot tighter, or suppose that suddenly your fleet size is reduced because of changes in chartering possibilities, or it has been pulled out from under your control and put in another agency temporarily. The flexibility or the cushion, the fat in the system, is suddenly carved out. Can we get in a position where the heuristics themselves start to fall apart?

SMALL: This ties in with what he was saying about complexity with respect to Figure 12.1, where you talk about doing tasks to see whether the heuristic is doing what it should be doing.

The obvious thing to do with a model is to run a small problem where you can, by some other means, come up with a solution, and compare your model against it.

Now, I am asking this as a question, because this is not particularly my area of knowledge. But, it sounds as if a heuristic might do very well on that model, in that situation. And then, when you increase the complexity and give it a bigger problem, it might in fact not be doing at all well, and you would never know. Is that accurate?

JARVIS: Yes, absolutely.

WEBSTER: In the things that I do, my heuristics involve decision making, so I incorporate some randomness in decision making and come out with a variety of different solutions, a large number of different solutions that have involved different decisions along the path. And then you can at least get a spectrum of answers. There may not be an optimal solution included, but you get a spectrum of answers from which you can select one that ranks well.

I think if you are careful not to exclude any possibilities in your heuristic, then, as long as you do things like run long enough or get enough examples, you can, in essence, come as close to optimality as you please.

DOENGES: On the choice considerations, I think a key thing there that can be tied to operator control is the understandability of the model. When you get a solution back and you see something that looks a bit strange, you want to know why did the model do what it did; you say "It does not seem right to me; I would have done something else." You really need to understand why the model put out what it did.

JARVIS: Put another way, "I will not let the model take control if I do not understand what it is doing."

DOENGES: We hope to never let the model take control.

JARVIS: I will not even accept an answer from a model if I do not understand it.

WEBSTER: The interactive business is essential in this kind of arrangement. A good interactive system permits you to make all of those decisions as a kind of input, and allows you to see the consequences of the way you have done it, so that you can make comparisons with whatever else it comes up with.

You can avoid this problem, if it looks strange, with good heuristics and with good interactive operating systems. You can put in what you think does not look strange, and see what the differences are, and perhaps learn something about your model or find something that is missing from your model--which is a likely possibility--or learn something subtle about what it is that you are assuming, and that it does have disastrous consequences.

It is a learning experience and a growing experience for both the model and for the user. That is important.

DOENGES: Maybe a potential problem is that some users cannot tolerate a learning experience. Maybe their needs are so critical that they have to have a good, understandable answer fast, and they cannot afford--either in terms of their time or whatever--a lot of sensitivity analysis.

DOUGLAS: It is one thing to have it fast, but it is another thing to be able to do it ahead of time, to find out whether you are going to like it or not. So, it does not have to be right now; if it is set up properly, you can do it ahead of time, so that you have a feel for what your reaction is going to be for a given set of circumstances.

DOENGES: Improve the technique. Gain confidence ahead of time.

WEBSTER: Two things are clear. One, you need a development period which is rather long, on any one of these things. And two, you need an intelligent operator. I do not think I have any hope for any of

these systems being capable of being used with confidence, without an intelligent operator.

Given these two, I think you will not reach the situation where you are so sensitive to the output.

TEMMLER: What do you consider a period that is very long?

WEBSTER: Years. A couple of years, depending on the complexities of the system. I think it would be unreasonable for us, for instance, to conceive of the SEASTRAT project being less than three years, say, of concerted effort. That would be my guess. You are talking about that kind of development.

TEMMLER: Or possibly the alternative. It may not be less than three different projects. You can fall back on the old saw-

DOENGES: You mean alternative approaches?

TEMMLER: No, three different phases or subsections. There is an old saw that a project over two years in length is not going to be manageable.

WEBSTER: We each have our own views on that. I am reminded of Paul Mentz's comment yesterday; you cannot get nine women to produce one baby in one month. I think that is an appropriate comment here.

TEMMLER: You do not produce an adult after nine months--and maybe they need the baby first, to nourish it into maturity.

WEBSTER: That is right. My experience is that these things require living with. They each have their own kind of personality, whether you are using an optimality approach or a heuristic approach. They have their quirks, their personalities. It requires an intelligent operator to develop the relationship with the system in order to use it properly. That is a time and spending thing.

DOUGLAS: Three years is not bad, as long as you have something in the meantime that can be used. I think that is very important. We cannot

just put something off for three years. As long as you are using part of it, then fine; take three years or five years.

KEYFAUVER: I would add a couple of things here. You mentioned three years. In the environment in which they are developing SEASTRAT, there are a couple of things that are going to happen in three years, one of which is that all of the military people who are involved in the development of this thing right now are going to disappear. They will be replaced by an entirely new set of people who do not know anything about it.

The other thing is that at the end of the project they will be the only ones who know anything about it, so they will end up with something they do not know anything about after that period, anyway.

MENTZ: The unfortunate part of what you say is that that will result in almost a guaranteed failure.

KEYFAUVER: Which is why most of the models that the government builds, one way or another, end up as failures.

MENTZ: And that is a sad commentary on the system. Unless that system is changed, we will not achieve what has been described as the objective here.

POLLOCK: Maybe we need to embed this into the educational process that will come into the shipping companies, and then the generation that comes in will just feel comfortable doing those kinds of things. I think that may be the way things are going to be done.

MENTZ: The comment was made relative to MSC. The companies may not share that same dilemma.

GALLAGHER: I would like to comment on that. The people who work in the ADP system for the Military Sealift Command and SEACOP have been there for 11 years. It has been the same people. The stability of the individuals within the organization will remain very high. I am the only military person in the whole organization that is doing the

ADP. That is not a very relevant comment.

But one that is relevant is, "how long does it take to develop a system?" We have been working with SEACOP, which is the heuristic model, for a long time. I have learned how to use it. It has a personality, and you begin to develop new and better ideas on how to apply it. But you do hit a point where you cannot patch it. You hit a point where you are at the end of the road. I mess with the heuristics in it so much that I do not know the impact when I mess with it. I may solve one little bitty piece and create 16 more problems further down toward the end of it.

It is time to get out of that and get ourselves a more up-to-date model, something that is more interactive with us at this point in time. That is what we are trying to do with it.

I am in total disagreement with the statement that it would disappear in three years. If I thought that, I could not do it. I want to make that point.

WEBSTER: One point you mentioned is a point of just strict data processing procedure and technique. Most of us do not redocument and clean up our programs as we change them. And we wind up with the same problem, and that is strictly an administrative procedure to insist upon. Unfortunately, that is not the standard that is insisted upon.

DOENGES: There is always something else that is more important; that is why.

Summary Presentations of
Discussion Group Leaders

SOLAND: For the first part of this afternoon we are going to have a series of reports, one by each of the six discussion group leaders. I have asked them to try to summarize and synthesize what went on in their respective sessions, where we attempted to look at six different aspects of this bundle of problems that we call "cargo ship routing and scheduling."

POLLOCK: I do not think it will take very long to summarize what happened in our session. I think the one thing that became clear is that the liner operators at this particular time are really not making great use of optimization techniques, but are really moving into what we would call heuristics, use of computers in simulation. We agree that there is some use that can be made of these techniques for the companies, but pure optimization is something that is not likely to be used: it is just something that the operators themselves have not really made much use of in terms of cargo ship scheduling and routing. However, they do use optimization techniques where they can be used, for instance, in scheduling of containers and things of that nature. That was one thing that came out of our discussion very clearly.

I think another point that really became clear is that there is probably a willingness, when problems are well defined, complex, and really beyond the capability of an individual to analyze, to use any technique that is available to solve the problem, whether it is heuristic or optimizing.

What I just said agrees, I think, with what Russ Stryker said in his talk yesterday, that basically the payoffs are not great enough. The problems are too complex. Uncertainty enters into it, so that in general we see the future going in the direction of on-line systems and more

simulation as being useful to help the operators in their problem areas. But generally, for pure optimization, there will not be a lot of use.

MACLEAN: We discussed three questions, including various aspects of them. The first question had to do with the state of the art in routing and scheduling. The second question dealt with the problems seen in current operations. The third one dealt with what work should be done to improve the capability for optimizing schedules.

It seems, on the basis of the first question, that we do not really have a definition of the problem as such. What is presently going on is that various segments of an operation have been allocated to the elements of the organization. There are manual processes, there are computer-assisted processes, and there are semiautomatic processes in various states of application. This varies according to the size of the organization and the complexity of the operation, as they see it.

A definition of the total process does not seem to be in place. In other words, taking the driving function and carrying it right on through to an objective function for a corporate policy does not seem to be exposed at the operating level, or at least it did not come out of what we discussed in our session.

It would appear that there are more important things to deal with in the operating model. Although it is suspected that there is money being wasted by not having a better idea of what should and can be done, there apparently is not the drive to do it, in the sense that the Airco illustration of yesterday would suggest it is going to be done.

The problems in the current operation are typically short term. They have to do with variations in the schedule, primarily, and they are not easily described. The algorithms that are presently being used deal with longer term problems, and the frequency with which they are being exercised may vary from once every few months to once every six months to a year. There is not any consistent pattern that has been identified as yet.

There are problems of identifying the characteristics of the operation. There are problems of establishing feedback with respect to post-operation analysis, so that you can determine whether what you thought was going to happen actually did happen, or whether, in fact, the performance is inadequate and better information can be fed into the next round of scheduling.

When it comes to identifying work that should be done with respect to improving capability for optimizing schedules, it seems quite clear that there has to be an objective function definition that is more than just making sure that the requirements immediately evident are being set and satisfied.

If we look at some of the particular examples, one can find that a refinery may be the driving function for the operation. It generates the demand for cargo movement, and as long as that demand is being adequately satisfied in the view of those who are placing the demand, there does not appear to be any particular drive to improve the operation. There is an operational window within which services should be made available and goods delivered. There does not appear to be any definition of penalty for early arrival, and unless things really collapse, there is not a well-defined penalty for late arrival.

In the absence of a clearly disastrous situation, attention does not seem to be applied to the improvement of scheduling. Although it appears that the driving function typically emanates from a marketing analysis and the development of corporate strategy, this does not filter down in any effective way into the operating area. Apparently the people who are involved in making these kinds of decisions are sufficiently remote from the actual operation that the relationship is not clear.

In this sense, there is a clear parallel between industrial operations and what we have heard with respect to the military operation. Whether the feedback avenues are adequate in any case is not clear, either,

because there apparently is not very much in the way of post-voyage analysis taken into account on any systematic basis, and there is no avenue by which this routinely goes back for a reassessment of effectiveness.

Those processes that have been computerized are apparently reasonably appreciated, but they are not typically used on a daily basis. It is not clear at this time that there is a particular thrust for doing that. I think that is about the essence of what we discussed.

BALLOU: In our group we discussed three areas of interest. As Captain Scott mentioned yesterday, there is, first, the course of action development in which you determine the feasibility of various courses of action as defined by commanders responsible for some military operation, given the constraints from a transportation point of view. In the second phase you flesh out one of those courses of action a little bit more, i.e., you give some more detail as to when certain units of interest are to arrive. And then, of course, there is the third phase, which is more akin to the standard scheduling problem.

The basic goal I set for the discussion in the group was to get some idea of what the community thought about the feasibility of attacking, and possibly solving, the problems as they were presented. A lot of the conversation, however, dealt with the inputs rather than model development, specifically in the cargo areas. We discussed such things as how much detail is available and the prioritization problem that was brought up yesterday.

To get to the bottom line, the general opinion I got from the group is that these problems are certainly feasible, at least at first sight.

One of the areas that I was particularly interested in, and it came up from the discussions, is the interaction between the various force-sizing phases. In a commercial sense, you are looking a year ahead, or six months ahead, as opposed to the actual scheduling: what is the link

between these two? If we make some commitments as to how we can move the various DoD assets around, they tend to get cast in concrete, and you cannot go back and say, "That was because the ships were positioned this particular way on the 5th of March." There has to be some link between a scheduling operation and the force-sizing one.

A suggestion was made that you cannot get away from the current type of force-sizing, that you should not go to complete scheduling when you do three-month or six-month planning. But it was suggested that perhaps you should go to almost a scheduling kind of routine when you start talking about your critical assets.

There was also some discussion as to what the objective function should be. Is cost a major factor? I think in some cases it might be, but it was also generally agreed that if you are fighting a war the cost may not be the number one criterion. And even then, cost was unsettling to some people, because instead of dollars, perhaps what it really should come down to is the management of the assets, the ships. Basically, they are your ships, and it is a national allocation of resources we are dealing with. Even though DoD may have overriding requirements for ships, you are not going to disrupt some POL pipelines. Thus, is a required delivery date a constraint or part of the objective function? I am not sure which one it goes into. There is the idea that you either meet it or you "fall on your sword," so to speak; I do not think we reached any resolution on that. It was discussed quite a bit. It is hard to quantify it unless you go back and talk to the shipper, or the consignee, the CINC, as to how important it is to get the cargo there by the RDD. This question was almost in the box of "too hot to handle."

To come to the bottom line on this thing, I was encouraged by the feedback I got from the group. The general consensus was that it is a problem that can be done. You can work all three areas. In the force-

sizing area the first part of the problem is probably the easier of the two. But none of it is beyond the realm of attack, and there is certainly hope for solution.

I would like to open it up now to members of the group that were there, to make sure that I covered the points and did not bias it in a direction away from the consensus of the group.

KASKIN: I think you should mention also that during the first fleet-sizing component, during the estimating that was discussed, that perhaps a solution would really be a set of parametric solutions. You would have several services versus cost that could be given to the CINC. For each, there would be a closure date and the resources that you would need to meet it. Under various assumptions you would develop and present various alternatives so that he could choose among them in his course-of-action-selection phase.

BALLOU: That is expanding on our answers to our bosses. The MSC, at least at the beginning stages, ought to look at groups of ships. Obviously, there are first the MSC-chartered ships. You look at berth term service. Can that solve the problem? Then you have voluntary charters, and then you get into the ready reserve forces, another group. In your first assessment of the problem, then, you really do not have to get down to exactly how many ships you need. You can look, at that point, at groups of ships.

GALLAGHER: Our group was just a little bit unique, I think, in what we expected and wanted out of this symposium. We were really looking for help to solve one of our immediate problems, one that I am involved in at this time. That is the problem I addressed in my group, of coming up with a new planning system called SEASTRAT. The exchange of information between myself, the military, DoD, industry, and the academics present was extremely important and was of great value.. It certainly gave us

three or four new sources to go to for some assistance, and I want to thank everybody for their fine inputs.

In our group I discussed SEACOP, our present system. It is very heuristic. I talked about how we have used it, how we have switched it around, how we have tried to improve it, and the reasons why we are going to develop our new system.

I think it became apparent early in the discussion that there really is not anyone who has dealt with as big a problem or as detailed a problem as we have in STRAT mode. We discussed various methods of attacking that problem, and someone even made the point that we are not going to solve anything until we get the problem properly identified. I wholeheartedly concur with that.

It was the consensus of the group that a mixture of solution techniques would be necessary to solve this problem; it is going to involve both heuristics and mathematical programming.

But whichever way we go, either with heuristics or mathematical programming, we are going to have to come up with some definite measures and tests, so that we will have some basis for checking the validity of whatever we are doing.

We talked about various ways to identify the problem, to formulate the problem, so that we could get it down to manageable size. Rather than looking at the whole problem with all of its variables, we would like to break it down. This would perhaps give us an opportunity to apply some of these techniques we have heard about, maybe not the same technique to the whole problem but different techniques to different subproblems that we have within this scheduling problem.

We talked about aggregating the problem to reduce its size. We talked about partitioning it, using decomposition, and perhaps breaking it down by time or geographic location.

Then we went into a somewhat lengthy discussion of the user interaction with a system. We felt that the human being, being what he is, would have to interact with the system, no matter what kind of solution we came up with. We felt there is always going to be a requirement for somebody to schedule the leftover stuff. There is always going to be a requirement for the user to get in there somewhere to influence the decision-making process. Exactly where that is going to take place, and when we are going to be able to do it, will depend a lot on how we get the problem formulated and our objectives established.

Of real interest to me was the developmental times that I heard expressed here by the various people who have their own systems, who have spent some time on what I would say are fairly narrow and well-defined problems. You are able to express such a problem, you are able to model the problem, and come up with a solution in terms of some objective, while dealing with the restricted constraints you had in there and a restricted number of variables. Your development times were still extensive.

I think the overall view here was that as we go down this road with SEASTRAT it will perhaps be a phased solution. We are not going to get it all at once; in fact, we are going to learn as we phase it in, and we are really going to be in a process of development as we do phase it in.

YOUNG: I think one other thing that we discussed was the objective function formulation and how it related to the RDDs, the delivery dates, in terms of those being rigid requirements versus putting priorities or weights on them and trying to meet some with higher priorities than others.

GALLAGHER: The point was made that an RDD of the 20th day for tanks is different than an RDD of the 50th day for ice cream; the present model treats them the same.

KASKIN: Were there any specifics mentioned on the development time and dollars?

GALLAGHER: There were various figures mentioned, depending on what you want to talk about. People were talking about anywhere from two to five years and more than one million dollars.

CHRISTOFIDES: We discussed at some length what is meant by exact methods, and we put down a number of points that have to be addressed before the execution of a model. How exactly does the model represent reality? What things are we allowed to do and still call the model exact? That is as far as the model is concerned, let alone the solution to the model. A number of points were brought up.

For example, it was mentioned that aggregated data are perhaps more reliable than unaggregated data, but nevertheless it is a simplification of the real problem. A number of solution procedures involve enumeration of routes, that is, partial enumeration; it was felt that perhaps this was not a very great limitation as far as the exactness of the model was concerned, that it was a perfectly acceptable thing to do. These are problems that are not related to uncertainty.

It was pointed out that, in general, uncertainties have to be incorporated in a real-world model. But this is not a problem for exact solution procedures only. Exactly the same problem arises with heuristic methods.

We tried to concentrate the discussion on what is specific to an exact solution method. It was pointed out that even with exact methods one should not insist on a proof of optimality, but rather know how close one is to an optimum solution, and provided that the guarantee is small enough one could call the method exact. A solution procedure in which you specify the error, and you can specify it to be small enough, should be considered exact. I think our discussion finished on an optimistic note on the part of most of the participants, who felt that would soon be the case. I personally think that is the state of the art now.

I think, in summary, it was felt by a number of people that over the next few years it will be possible to exactly solve average-sized ship routing and scheduling problems.

KASKIN: Do you think that the problem has been sufficiently well defined for you to have that level of optimism? Do you think the people here have an adequate feeling for what we are talking about?

CHRISTOFIDES: I am sure a lot of the people here have a good enough feeling for what the problem is. I personally think it has not been defined succinctly enough. I would like to see a more strict definition of some core problem, not necessarily a problem that is a real-world problem, but one that, if solvable, could be applied almost immediately to a real-world problem. I think this is something that has to be done over the short term if progress is to go on.

JAIKUMAR: Some of the problems seem to be fairly well defined. I think we can have exact solutions in the next two to four years.

CHRISTOFIDES: I agree with that. It is simply that a lot of different types of problems have been mentioned here, and in trying to extract from them a set of central problems, I do not feel confident that I can identify a number of them. I think that I, perhaps, have prejudices, since I know the road transport system better than the shipping systems; I can better identify the problems that have a closer relation to that former area.

It was pointed out during the meeting that although with heuristics you can deal with problems that are fuzzy and ill-defined, as real problems are, simply because the heuristics themselves are fuzzy, one does not have this ability with exact methods. In order to apply an exact method you need a very well defined problem. I think there is some work necessary in trying to identify some core problems that are significant problems in ship routing, which one can then go and work on. Because of the success of some of the methods that we are using in similar problems, I am confident that it is possible to exactly solve such problems, of moderate size, in the pure form.

SOLAND: I would like to expand upon that, in the sense that I think we felt that exact or almost exact solutions could be found for well-formulated

problems that modeled the real-world problems, or at least are core models for the real-world problems.

My guess is that the academic people here who have worked with the mathematical techniques have perhaps seen enough to feel that they could come up with core problem models, and then could go on to find exact or near-exact solutions to them. They are not prepared to go home and immediately write down a formulation of such a core model, however. But if they were given the opportunity to spend a reasonable amount of time working with the people who have the problem, and interacting, then they could come up with core models that they could then work on the solutions for. We are at an intermediate point.

JARVIS: There was a great deal of discussion about the role of heuristics and the development of heuristics, the kind of heuristics that might be produced in the future, and the current state of heuristics. We looked at a breakdown that was suggested, one that went all the way from those that have rule-of-thumb estimates and procedures that give you a little better feel for the optimality values, all the way to what we called optimality-based heuristics, those that are derived from optimality models of the system which we cannot solve.

We were very careful to try to define what we meant by "heuristics." I may have influenced the group's discussion there. There was also some give-and-take on the actual modeling process itself, from the realistic view that sometimes you simply cannot solve the problem, so there is no sense in writing it down in any exact form. So we looked at the broad spectrum of heuristic approaches, which is backing off on the model itself as well as possibly backing off on the solution techniques and the techniques for analysis.

We also looked at another categorization of methods which went all the way from manual methods to interactive methods, which would include the human and the computer inside the loop, all the way to fully automatic methods. And I guess our conclusion was that we would like to move away from manual methods. A number of people are doing that, some with more success than others.

By and large, we do not believe that we will ever achieve fully automatic methods in any realistic time frame. In fact, there was some question as to whether or not we even would want to have fully automatic methods that would be completely beyond operator control, i.e., just produce a solution which you would blindly accept. Although there is some question about what is meant there with regard to that approach, there seemed to be some heated discussion in the group as to whether or not more or less modeling ought to take place, and just what degree of modeling should be involved in the process.

I think there was a general sense that we ought to try to do a better job of utilizing at least suboptimization in some of the analyses that we are doing, but that we should not pay the price, on a cost-effectiveness basis, that would slow down the development of the techniques or hinder the application to real problems.

We had some discussion on quality characteristics and quality evaluation of heuristics. We recognize that there has not been much done in terms of theoretical analysis that would really give us some feeling of security in choosing a particular technique.

We think that more could be done on experimental analyses that would help with some cross-comparison between operating groups. Possibly a better job of communication can be done in terms of establishing a set of test problems, as has been done in the case of the vehicle routing problem. Also, there could be more comparison of different models on the same set of data to get some kind of evaluation of where we are relative to our ability to manage problem environments.

A smaller amount of discussion was concentrated on problem complexity. There was a feeling conveyed that heuristics are somewhat robust and independent of problem size, although there was another concern that as the availability of resources became tighter and tighter, for example, in closing down time windows or reducing fleet size, that

possibly feasibility would be harder to obtain, and this might cause the heuristics to have to work harder and harder, and eventually cause you to change their nature.

And finally, with regard to choice considerations of how to select a heuristic, we identified several possible criteria, not the least of which would be availability, as well as ability to react, capability for operator control, and certainly quality. One final one that was added to the list was understandability. We would like to make sure that what we are developing and what we are using, whichever side we are on, is at least understandable to the user and understandable to the operator and the scheduler.

I guess we feel that there is a good climate for heuristics, and they are certainly not going to go away. We would like to continue their encouragement and move toward greater and greater interface and capability for modeling and analysis.

KASKIN: Was there any determination of which of the problems presented would be most amenable to heuristic solution? Or was the discussion a general one for all of the scheduling problems that were discussed?

JARVIS: I do not think we got down to specific problems. I guess the general feeling was that the problems we were looking at, with regard to temporal considerations, ports, fleet sizes, commodities, time windows, and so forth, were just far too complex to admit exact solution procedures, so we did not really focus in on the particular problems that would be amenable. We just felt that, generally, heuristics were going to be required.

CALLAGHER: Our model, SEACOP, is heuristic. It does work, but it takes a lot of effort. Our objective here was to look around and see what was out there, whether there was anything to assist us as we proceed in the development of SEASTRAT. My opinion is it is great to have the academic people continue to charge along in the exact solution mode. Unfortunately, I do not have five or ten years to come up with an adequate solution.

As I look at my problem, as it has been defined, there must be at least a thousand solutions to it, and I certainly do not really care to get an optimal one at this time. I want one of those thousand to be reasonable so I can use it. It would appear to me that heuristics, with their ability to break the problem into smaller pieces, provide the practical approach I need.

MACLEAN: One of the things that came out in our discussion is that a large number of the problems that are faced in scheduling are simple vessel assignment problems, and by tackling this to get the major portion of activity underway, one then has time to address the more difficult aspects of the scheduling problem. Part of your solution may be to get this 60 to 80 percent or so of your problem out of the way, and then focus your attention on the more difficult aspects of it.

JARVIS: I do not think we focused our attention on some of the submodels nor did we intend the whole procedure to be heuristic in its elements. Certainly, if one of the elements is an assignment problem, and somebody can work on that aspect of it to get an exact technique for the assignment problem, that is great. I do not think that anybody wants to discourage that kind of development. We were looking at the broader, overall problems, and in that regard what needs to be done to bring along the total integration of all of the submodels into some operating environment.

GALLAGHER: That was the consensus of our group. We felt it is going to be a hybrid approach. It looks like a heuristic for the first cut, and then when you get down to suboptimizing, we might want to get into exact solutions to particular subproblems.

WILMES: Tom Baker's presentation is the kind of thing that we are looking for. What it does in business, it could do for us, at least as a first cut.

Final Discussion Prior to
the Close of the Symposium

SOLAND: We will keep this final discussion fairly short, but I think we can give people a chance for a wrap-up of some sort. Perhaps I can pose a couple of meaningful questions. I think from the discussion group leaders' reports we have perhaps gotten a fairly good idea of what we have learned and where we are with respect to quantitative models in cargo ship routing and scheduling.

Let me ask these gentlemen up here; have we learned enough about the criteria to be used in the quantitative models? Have we really identified them?

MACLEAN: I think we have not learned enough about the criteria. Nobody seems to be prepared to quantify them, much less explicitly define them. And I think that is a serious problem. If you look at a model, somehow you have to recognize that in the commercial sector, e.g., sales drives production of goods and services. The production drives the need for transport services through cargo requirements. The criteria as to how good or bad the operation is with respect to the scheduling of these services depends upon who is identifying and setting these criteria.

Right now it does not appear that there is any formal procedure for doing this that is well recognized. Each individual organization seems to have its own procedure. And it is segmented in such a fashion that you do not get an overall view of this, and consequently it is sort of an individualized action that is involved. And I think that is a real problem.

We need to be more specific in identifying what the criteria are in arriving at some reasonable and selective means for quantifying them.

KASKIN: Some models have been built. Exxon and Bethlehem Steel, for example, have developed criteria to determine whether they are doing better now than they did before. It seems that some of these criteria can be developed, even

though they may be unique in each particular setting. I do not think it is a problem that cannot be solved.

SOLAND: I might ask, where do we go from here? We have had some indications of that, but does anybody else want to perhaps repeat or else just suggest something, or project something?

KEYFAUVER: I would like to say one thing. Whatever happens with MSC's model, they are planning to develop it. This problem is not going to go away any time in the future. I would like to see somebody, ONR or somebody else, put some money into long-term research that we need to find better techniques to solve these problems.

SOLAND: To carry that a step further, while it may not be feasible for the Department of Defense agencies to allow three to five years to develop an implementable system because they have got to get things up and going faster than that, there does seem to be real value in somewhat longer term support of directed research, in terms of modeling and solution methods, that will perhaps in five to ten years produce workable systems that can handle some of the problems that we perceive today as being too large.

KASKIN: I certainly agree with what Carroll Keyfauver has to say. And one of the recommendations I make when I go back to MSC Research will be to see if we can establish such a relationship with ONR. Research has been sponsored in the past through this mechanism.

Tony English mentioned previously that the Office of the Secretary of Defense is also interested in these large problems. Is there any formalized program in place, or going to be established, to possibly support development of such modeling techniques?

ENGLISH: I am new over there, but to my knowledge the answers are "no" and "maybe."

MACLEAN: One of the things that strikes me about this area is that you really do not know what problems you need to work on until somebody that

owns the requirements effectively stands up and says, "I have a problem." The researcher doesn't really know. There has to be some excitement for the researcher to say, "That sounds like an interesting problem. I did something that might apply. Let's see what I can do."

I think there needs to be some fundamental support in this area to excite researchers to formulate solutions to problems that seem amenable to solution. One of the things I think about the mathematics world is that the so-called "inverse solution method" is very valid. And one of the reasons that we have as much power, mathematically, as we have today is because in past decades a lot of mathematicians played--they formulated solutions, and tried to fit them to problems, and out of this you build up a catalogue of capabilities.

I would suggest that in some respects that is not the kind of position we are in right now. We do not have a very large catalogue of problems with solutions from a theoretical point of view. I would think that we need to excite the researchers to some play. It is not a very efficient way, but in the long run it may be very effective.

KASKIN: Do you think that MARAD can participate in funding some of this "playing" in the near future?

MACLEAN: You would have to talk to somebody else. I will tell you right now that anybody who reads the newspapers knows that MARAD's budget has been cut and cut and cut, so that today we are operating on an R&D budget that is half of what was submitted a year ago and about one-third of what it was ten years ago. We have thus been squeezed so heavily in this area that I do not see that there is any reasonable way to do these kinds of exercises within our budget. If it is going to become an effort of consequence, there is going to have to be a thrust in that direction that puts some resources there. They simply are not there now.

NEVEL: The industry itself can take up some of the slack, I think, if it can see the payback or benefit from this work, especially that from some

of the new techniques. And in listening to the conversations, it seems there is definitely application to what we call the tramp segment, however you want to define the word "tramp." I think the merchant marine is going to experience a change in the next few years, a continuing change that we have already seen to some extent. You are going to have fewer ships, but faster and more productive ships. It is going to be up to each operator to maximize efficiency or the payback on his investment, any criterion you want to use, to make sure that his operation is survivable and successful.

I think that the operators are going to be forced to take another look at some of these techniques, to make sure that all of the bases are covered, so that they can meet the challenge of the times and the challenge of the competition.

I think there will be a transfer effect of these payoffs to military applications. Remember, during wartime you always talk about the military taking over a large part of the merchant marine, but in fact we would have many operators operating fleets in support of DoD requirements.

SOLAND: I would like to come back to Walt Maclean's comments. First of all, I think there are already a certain number of key model types that academics work on, play with, and try to improve upon. Some examples are the traveling salesman problem, the quadratic assignment problem, other types of assignment problems, certain routing problems, and shortest path problems.

I think a very valuable part of this symposium has been to expose a number of academic researchers to some of the problems in ship routing and scheduling. If they can get some support and more explicit information about these problems, they are going to work on them because they are interesting enough, in terms of application and importance, and you will find that progress will be made.

BALLOU: I wonder if one of the operators would care to comment on Bob Nevel's remark. Do the ship operators see enough there to do work in this area?

BISHOP: With the possible exception of EXXON, which seems to do research for the sake of research more than anybody else, I do not see the international oil companies really working on these problems in a big hurry.

There is an oversupply of tonnage in all markets. A ship's time is not worth very much right now. Furthermore, it is very difficult to quantify what is being proposed. What am I saving in days and dollars? It must be presented in some way that management will buy it. I think management, at least in our company, does believe there is something there. They are willing to pursue it.

I have money in my budget, which, incidentally, was cut in half this year, to pursue research on a lower scale. In other words, we are not backing out. We just want to move along slowly, and see what is there. Maybe something will come out of it.

But what I conclude from some of these discussions is that solutions to our problems are not necessarily going to solve MSC problems. Our problems are smaller, and somewhat different, though some of them may very well be similar. But, particularly in the case of your long-range planning, I do not see industry coming up with methods and models that will help.

CHRISTOFIDES: Don't you think there is some benefit to be had by sitting down with OR people to crystallize the problems you are mentioning?

BISHOP: We are doing it right now, with both in-house people and outside contractors. We are just not pursuing it with full manpower and at full speed. We are not letting it die.

I think we are going to continue, but if you are looking for industry to provide in three to five years a rapid new scheme to solve your problems, I do not think it will come to pass.

DOUGLAS: I do not think MSC's problems are going to be solved by some sort of "magic", by somebody "playing" and seeing the perfect application for some particular little nugget. We have discussed the potential for a combination of approaches, and as long as nobody expects some big magic, I

think progress can be made. MSC's problem is unique enough that it is going to take just work on that problem, with the flexibility in mind of how to go at it, to come up with a solution package.

SOLAND: Are there any other comments? I would like to thank you all very much for your participation in this symposium. I think we have all gained something and we have all contributed something.

REFERENCES ON CARGO SHIP ROUTING AND SCHEDULING

(Compiled by David Ronen)

1. Almogy, Y., and O. Levin (1970). Parametric analysis of a multi-stage stochastic shipping problem, IFORS, OR 69: Proceedings of Fifth International Conference on OR, Tavistock, London, pp. 359-370.
2. Appelgren, L. H. (1971). Integer programming methods for a vessel scheduling problem, Transportation Science, 5, pp. 64-78.
3. Appelgren, L. H. (1969). A column generation algorithm for a ship scheduling problem, Transportation Science, 3, pp. 53-68.
4. Baker, T. E. (1981). Interactive vessel scheduling at Exxon, Presented at CORS/TIMS/ORSA Joint National Meeting, Toronto.
5. Baumol, W. J., and P. Wolfe (1958). A warehouse location problem, Operations Research, 6, pp. 252-263.
6. Bellmore, M., G. Bennington, and S. Lubore (1971). A multivehicle tanker scheduling problem, Transportation Science, 5, pp. 36-47.
7. Berten, B. (1973). Mathematical description of stochastic processes for dispatching ships, Seewirtschaft, 5, pp. 99-101 (in German).
9. Bodin, L., B. Golden, A. Assad, and M. Ball, (1981). The state of the art in the routing and scheduling of vehicles and crews, College of Business and Management, University of Maryland, College Park, MD, Working Paper MS/S #81-035.
9. Boffey, T. B., E. D. Edmond, A. I. Hinxman, and C. J. Pursglove (1979). Two approaches to scheduling container ships with an application to the North Atlantic Route, Journal of the Operational Research Society, 30, pp. 413-425.
10. Bratley, P. and M. Florian (1972). Scheduling newsprint deliveries made by tugs and barges, Proceedings of NATO Conference on the Application of Operations Research to Transport Problems, Western Periodicals Co., North Hollywood, CA, pp. 203-214.
11. Briskin, L. E. (1966). Selecting delivery dates in the tanker scheduling problem, Management Science, 12, pp. B224-B233
12. Cheshire, I. M. (1972). Most profitable scheduling of bulk cargo fleets, Norwegian Shipping News, No. 11C, pp. 40-41
13. Clarke, G. and J. W. Wright (1964). Scheduling of vehicles from a central depot to a number of delivery points, Operations Research, 8, pp. 568-581.

14. Clausen, C. O. (1981). Vehicle routing algorithms for local delivery at naval supply centers, Unpublished Master thesis, Naval Postgraduate School, Monterey, CA.
15. Conley, J. H., R. S. Farnsworth, E. Koenigsberg, and U. Wiersema (1968). A linear programming approach to the total movement of a homogeneous product, Transportation Science, 2, pp. 289-302.
16. Dantzig, G. B. and D. R. Fulkerson (1954). Minimizing the number of tankers to meet a fixed schedule, Naval Research Logistics Quarterly, 1, pp. 217-222.
17. Dantzig, G. B., W. O. Blattner, and M. R. Rao (1967). Finding a cycle in a graph with minimum cost to time ratio with application to a ship routing problem, in Theory of Graphs, (P. Rosenstiehl, ed.) Gordon and Beach, New York, pp. 77-83.
18. Datz, I. M. (1968). Planning tools for ocean transportation - ship scheduling, Norwegian Shipping News, pp. 1064-1069.
19. Flood, M. F. (1954). Application of transportation theory to scheduling a military tanker fleet, Operations Research, 2, pp. 150-162.
20. Hayman, C. (1972). Fleet scheduling - best by computer, Seatrade, (April) pp. 65-66.
21. Hunt, R. B. and E. F. Rosholdt (1960a). The concepts of notional ship and notional value in logistics capability studies involving merchant ships. Naval Research Logistics Quarterly, 7, pp. 1-6.
22. Hunt, R. B. and E. F. Rosholdt (1960b). Determining merchant shipping requirements in integrated military planning. Naval Research Logistics Quarterly, 7, pp. 545-575.
23. Johannson, K. (1969). A heuristic approach to the multi-discharge routing problem, Unpublished Ph.D. Dissertation, Illinois Institute of Technology.
24. Koenigsberg, E. and R. D. Lam (1976). Cyclic queue models of fleet operations, Operations Research, 24, pp. 516-529.
25. Kydland, F. (1969). Simulation of Liner Operations, Institute for Shipping Research, Bergen.
26. Laderman, J., L. Gleberman, and J. F. Egan (1966). Vessel allocation by linear programming, Naval Research Logistics Quarterly, 13, pp. 315-320.
27. Lawrence, S. A. (1972). International Sea Transport: The Years Ahead, Lexington Books, Lexington, MA.

28. Levy, V. D., S. P. Lvov, and S. E. Lovetsky (1977). Man-machine system for merchant fleet operation scheduling. Proceedings, 7th International Symposium on Transportation and Traffic Theory, Kyoto. (Institute of Systems Science Research), pp. 823-836.
29. Mathis, S. J. Jr., (1972). Synthesis of a crude oil supply system: A non-convex network problem, Unpublished Ph.D. Dissertation, Case Western Reserve University.
30. McKay, M. D. and H. O. Hartley (1974). Computerized scheduling of seagoing tankers, Naval Research Logistics Quarterly, 21, pp. 255-264.
31. McKenzie, J. S. (1971). The routing of ships, Journal of Transportation Economics and Policy, pp. 201-215.
32. Nemhauser, G. L. and P. L. Yu (1972). A problem in bulk service scheduling, Operations Research, 20, pp. 813-819.
33. Neuhof, B. (1974). The ship line decision problem on ports of call, Hansa, 23, pp. 2013-2017 (in German).
34. Norman, R. J. (1973). An algorithm for the scheduling of vessels through the Panama Canal, Unpublished Master thesis, Naval Postgraduate School, Monterey, CA.
35. Naslund, B. (1970). Combined sea and land transportation, Operational Research Quarterly, 21, pp. 47-59.
36. O'Brien, G. G. and R. R. Crane (1959). The scheduling of a barge line, Operations Research, 7 pp. 561-570.
37. Olson, C. A., E. E. Sorenson, and W. J. Sullivan (1969). Medium range scheduling for freighter fleet, Operations Research, 17, pp. 565-582.
38. Pruzan, P. M., and J. T. R. Jackson (1967). The many product cargo loading problem, Naval Research Logistics Quarterly, 14, pp. 381-390.
39. Rao, M. R. and S. Zionts (1968). Allocation of transportation units to alternative trips--a column generation scheme with out-of-kilter subproblems, Operations Research, 16, pp. 52-63.
40. Ronen, D. (1979). Scheduling of vessels for shipment of bulk and semi-bulk commodities originating in a single area, Unpublished Ph.D. Dissertation, The Ohio State University.

41. Schechter, J. (1976). A minimum cost vessel routing model applied to the U.S. Trust Territory of the Pacific Islands, Maritime Studies and Management, 3, pp. 195-204.
42. Stott, K. L.. Jr. and B. W. Douglas (1981). A model-based decision support system for planning and scheduling ocean-borne transportation, Interfaces, 11(4), pp. 1-10.
43. Schwartz, N. L. (1968). Discrete programs for moving known cargoes, from origins to destinations on time at minimum bargeline fleet cost, Transportation Science, 2, pp. 134-145.
44. Whiton, J. C. (1967). Some constraints on shipping in linear programming models, Naval Research Logistics Quarterly, 14, pp. 257-260.

SYMPORIUM PROGRAM

FEBRUARY 3, MORNING

Opening Remarks, Harold Liebowitz, Dean,
School of Engineering and Applied Science, and
Richard M. Soland, Department of Operations
Research, The George Washington University.

Department of Defense Ocean Transportation Routing
and Scheduling Requirements, John H. Scott, Military
Sealift Command.

Characteristics of Commercial Ocean Carrier Routing
and Scheduling Problems, Russell F. Stryker,
Maritime Administration.

A Review of Cargo Ship Routing and Scheduling Models,
David Ronen, University of Missouri-St. Louis.

FEBRUARY 3, AFTERNOON

The Utility of Heuristics in Relatively Complex
Strategic Mobility Models, Carroll J. Keyfauver,
General Research Corporation.

Mathematical Programming Models in Cargo Ship
Routing and Scheduling: Theory and Computational
Results, Marshall Fisher, University of Pennsylvania
and Ramchandran Jaikumar, Harvard University.
(not appearing in these proceedings)

Panel Discussion

FEBRUARY 4, MORNING

Introductory Remarks Concerning Discussion Groups,
Richard M. Soland.

Meetings of Discussion Groups 1 and 2.

Meetings of Discussion Groups 3 and 4.

Meetings of Discussion Groups 5 and 6.

FEBRUARY 4, AFTERNOON

Interactive Vessel Scheduling at EXXON,
Thomas E. Baker, EXXON Corporation

Presentations by Discussion Group leaders

Panel Discussion

Closing Remarks, Richard M. Soland

Discussion Groups and Leaders

1. Industrial Operations - Walter M. MacLean
2. Liner Operations - Donald K. Pollock
3. Crisis Environment/Time Sensitive Scheduling -
Joseph F. Ballou
4. Contingency Planning and Scheduling - Lawrence A. Gallagher
5. Exact Solution Approaches - Nicos Christofides
6. Heuristic Solution Techniques - John J. Jarvis

LIST OF PARTICIPANTS

Arjang A. Assad
University of Maryland
College Park, MD

Edward K. Baker
University of Miami
Coral Gables, FL

Thomas E. Baker
EXXON Corporation
Florham Park, NJ

Michael O. Ball
University of Maryland
College Park, MD

Joseph F. Ballou
Military Sealift Command
Washington, DC

Gregory S. Bentley
The Yardley Group, Inc.
Philadelphia, PA

Bruce R. Bishop
Standard Oil Company of California
San Francisco, CA

Lawrence D. Bodin
University of Maryland
College Park, MD

William V. Brierre, Jr.
Lykes Brothers Steamship Company
Washington, DC

Darryl J. Bullock
Science Applications, Inc.
McLean, VA

Nicos Christofides
Imperial College
London, England

James G. Conway
United States Lines
Cranford, NJ

Richard J. Coppins
Virginia Commonwealth University
Richmond, VA

G. Robert Doenges, Jr.
Science Applications, Inc.
McLean, VA

Burnie W. Douglas
Bethlehem Steel Corporation
Sparrows Point, MD

Charles E. Emberger
David Taylor Naval Ship Research and
Development Center
Bethesda, MD

J. Anthony English
Office of the Secretary of Defense
Washington, DC

John S. Foard
Organization of the Joint Chiefs
of Staff
Washington, DC

Terry L. Friesz
University of Pennsylvania
Philadelphia, PA

Lawrence A. Gallagher
Military Sealift Command
Washington, DC

Robert S. Garfinkel
University of Tennessee
Knoxville, TN

Neal D. Glassman
Office of Naval Research
Arlington, VA

Robert E. Greer, Jr.
Military Sealift Command
Washington, DC

LIST OF PARTICIPANTS-continued

S. W. Hall, Jr.
Joint Deployment Agency
MacDill AFB, FL

Ramchandran Jaikumar
Harvard Business School
Boston, MA

Chester J. Jakowski, Jr.
Military Sealift Command
Washington, DC

John J. Jarvis
Georgia Institute of Technology
Atlanta, GA

David G. Jesmer
Organization of the Joint Chiefs
of Staff
Washington, DC

James L. Johnson
Office of the Secretary of Defense
Washington, DC

Jonathan D. Kaskin
Military Sealift Command
Washington, DC

Carroll J. Keyfauver
General Research Corporation
McLean, VA

Walter M. Maclean
National Maritime Research Center
Kings Point, NY

Roy E. Marsten
University of Arizona
Tucson, AZ

James Mays
Oceanroutes, Inc.
Palo Alto, CA

Alan W. McMasters
Naval Postgraduate School
Monterey, CA

Paul B. Mentz
Maritime Administration
Washington, DC

Robert O. Nevel
Maritime Administration
Washington, DC

Eugene K. Pentimonti
American President Lines
Oakland, CA

Donald K. Pollock
Maritime Administration
Washington, DC

Richard F. Powers
Insight, Inc.
Alexandria, VA

Harilaos N. Psaraftis
Massachusetts Institute of Technology
Cambridge, MA

David Ronen
University of Missouri-St. Louis
St. Louis, MO

Linus Schrage
University of Chicago
Chicago, IL

John E. Scott
Military Sealift Command
Washington, DC

Duane W. Small
Science Applications, Inc.
McLean, VA

Richard M. Soland
The George Washington University
Washington, DC

Russell F. Stryker
Maritime Administration
Washington, DC

LIST OF PARTICIPANTS-continued

Robert D. Temmler
Moore McCormack Lines
New York, NY

William C. Webster
University of California-Berkeley
Berkeley, CA

John J. Wilmes
Military Sealift Command
Washington, DC

Stephen K. Young
Science Applications, Inc.
McLean, VA

2 - 8

DT